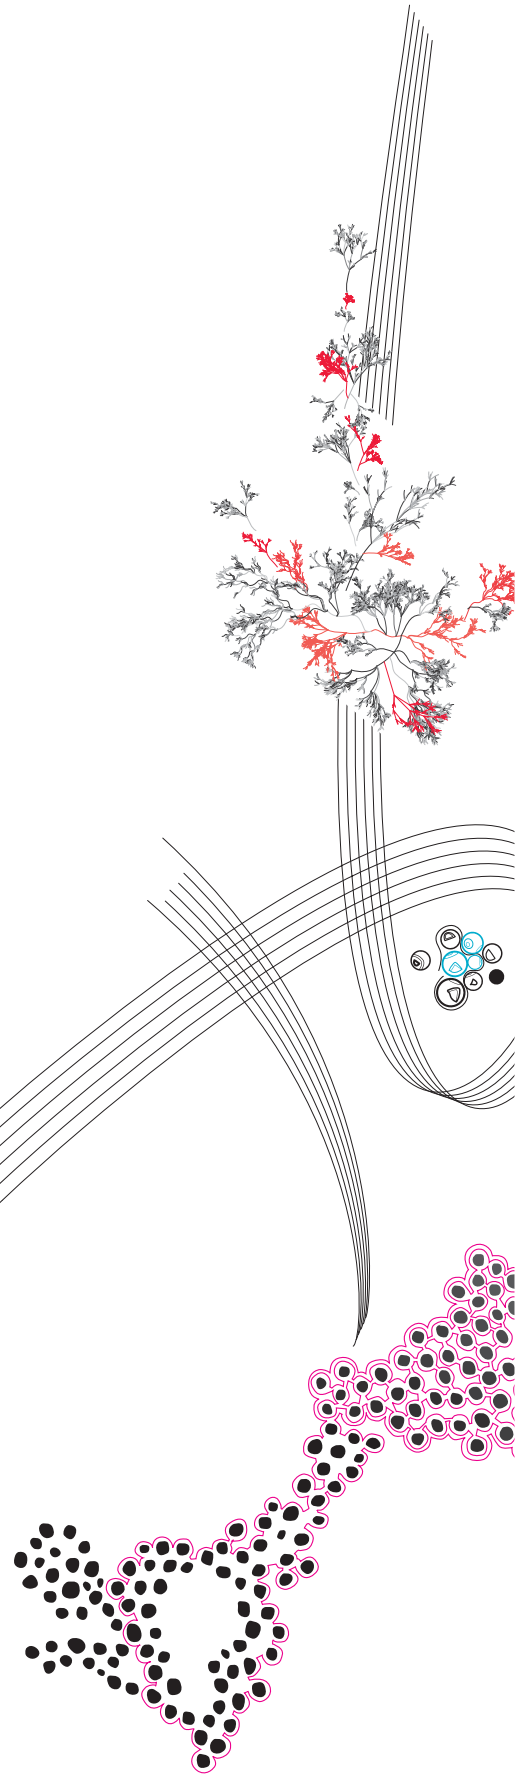

Computer Science and
Biomedical Engineering
Master's thesis

Quantified Photoacoustics for Carotid Artery Imaging

Max Rietberg

Graduation committee:
Dr. Francis Kalloor Joseph
Dr. Nicola Strisciuglio
Dr. Jelmer Wolterink
Dr. Estefanía Talavera Martínez
Prof.Dr.Ir. Wiendelt Steenbergen

December, 2023

Biomedical Engineering:
Faculty of Science and Technology

Computer Science:
Faculty of Electrical Engineering,
Mathematics and Computer Science

University of Twente

UNIVERSITY OF TWENTE.

Contents

Abstract

Since the very beginning, humanity has asked itself some fundamental questions: who are we? what is our purpose in life? is there life after death? Unable to answer any of these, in this thesis I will introduce and demonstrate a novel method to quantify photoacoustic imaging. Annually, vascular diseases worldwide result in 17.8 million deaths, including ischemic stroke (23%) caused by unstable thrombogenic atherosclerotic vascular wall lesions in the carotid artery. The decision of intervention for this condition is currently based on the degree of narrowing of the artery (stenosis), while plaque composition and associated vulnerability are more reliable predictors of future stroke events. However, no single imaging modality has been able to accurately identify vulnerable plaques. Recently, photoacoustic imaging has emerged as a promising medical imaging modality due to its ability to combine high light contrast with ultrasonic imaging depth. It offers submillimeter resolution and an imaging depth of several centimetres, which is sufficient for carotid plaque imaging. However, quantifying plaque composition is challenging as it requires knowledge of complex light propagation in tissue. To address this, I will showcase the utilisation of a reference chromophore, such as arterial blood, as a marker for light distribution (fluence) in simulation. In order to do this, I implemented the concept of a reference chromophore into three separate techniques: two baseline methods, and one using a modified version of NIRFAST. The 3 different methods were compared, and the behaviour of the virtual detector method was characterised. Finally, I demonstrated the accuracy of the novel method on a tissue-mimicking phantom. It was found that in their current forms, the techniques with a reference chromophore as a prior do not outperform those without. However, the characterisation of the proposed method shows its potential so that with appropriate modifications it could be used in plaque characterisation.

Samenvatting

Sinds het begin der tijden heeft de mensheid zichzelf een aantal fundamentele vragen gesteld: wie zijn wij? wat is het doel in ons leven? is er leven na de dood? Omdat ik geen van deze vragen kan beantwoorden, zal ik in dit verslag een nieuwe methode om foto-akoestische beeldvorming te kwantificeren introduceren en demonstreren. Jaarlijks leiden vaatziekten wereldwijd tot 17,8 miljoen sterfgevallen, waaronder ischemische beroerte (23%) veroorzaakt door instabiele trombogene atherosclerotische vaatwandlaesies in de halsslagader. Op het moment wordt de beslissing voor het ingrijpen bij deze aandoening gebaseerd op de mate van vernauwing van de slagader (stenose), terwijl de chemische samenstelling van de plaque en de daarmee samenhangende stabiliteit beter de kans op beroertes kan voorspellen. Echter, geen enkele beeldvormingsmodaliteit kan tot nu toe kwetsbare plaques nauwkeurig identificeren. Onlangs is foto-akoestische beeldvorming naar voren gekomen als een veelbelovende medische beeldvormingsmodaliteit vanwege het vermogen om het hoge contract van licht te combineren met de beelddiepte van ultrasoon geluid. Het biedt een submillimeter resolutie en een beelddiepte van enkele centimeters, wat voldoende is voor beeldvorming van carotisplaques. Het kwantificeren van de plaquesamenstelling is echter een uitdaging, omdat hiervoor kennis nodig is van de complexe lichtvoortplanting in weefsel. Om dit aan te pakken, zal ik in dit verslag laten zien dat een referentiechromofoor, zoals arterieel bloed, als een marker voor lichtverdeling (fluentie) kan dienen in simulatie. Hiervoor heb ik het concept van een referentiechromofoor in drie afzonderlijke technieken geïmplementeerd: twee basismethoden en één met behulp van een aangepaste versie van NIRFAST. De 3 verschillende methoden zijn vergeleken, en het gedrag van de virtual detector method is gekarakteriseerd. Ten slotte demonstreerde ik de nauwkeurigheid van de nieuwe methode op een weefsel-nabootsend fantoom. Op basis van dit onderzoek is duidelijk geworden dat dat de technieken met een referentiechromofoor in hun huidige vorm niet beter presteren dan de technieken zonder. De karakterisering van de voorgestelde methode toont echter het potentieel ervan aan, waardoor deze met de juiste aanpassingen kan worden gebruikt bij de karakterisering van plaques.

Chapter 1

Introduction

1.1 Clinical context

One of the leading causes of mortality worldwide is stroke, of which around 80% is classified as acute ischemic stroke [1], where the blood supply to the brain is reduced or interrupted. A large degree of ischemic strokes can be traced back to carotid artery disease [2], where (unstable) plaques narrow this artery. Unstable plaques should be removed surgically, with a so-called carotid endarterectomy [1]. Determining the risk of plaque disruption is a crucial but challenging task. Overestimating this risk may lead to unnecessary medical procedures or treatments, which can cause potential harm to the patient and increase healthcare costs. However, underestimating the risk of plaque disruption may result in a delayed or missed diagnosis, possibly resulting in serious harm to the patient, including permanent damage to the brain, and in some cases, death. Current methods of determining plaque stability rely, among other things, on assessing the degree of carotid stenosis (see Figure 1.1), ideally through non-invasive imaging modalities [3]. Modalities for determining plaque dimensions and composition like CT, MRI and nuclear imaging studies all have their disadvantages, mostly because they are relatively expensive and/or use ionising radiation. This is where photoacoustics imaging shows potential.

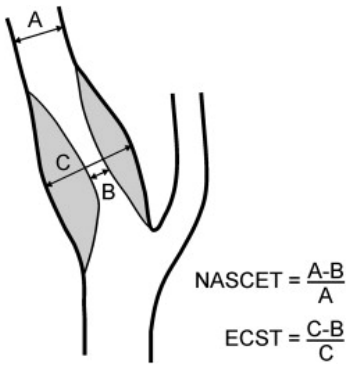

FIGURE 1.1: Two methods that are currently used to assess the degree of carotid stenosis, reproduced from Oates *et al.* [4]

1.2 Technical context

For medical imaging, a cheap, non-invasive, non-ionising imaging modality with high spatial resolution, high sensitivity, and good specificity is the holy grail. Photoacoustics (PA) is a relatively new modality that combines the benefits of optical and the benefits of acoustic techniques [5], which allows it to achieve high sensitivity and good specificity. In PA, a sample is (usually) illuminated by a source of non-ionising visible or near-infrared pulsed light. These transmitted photons are absorbed by various chromophores that occur naturally in tissue, raising them to an excited state. These excited molecules will convert the energy into heat, assuming that the thermalisation (non-radiative decay) dominates over radiative decay. This results in a small rise in local temperature and a connected rise in the local pressure. This initial pressure will propagate throughout the sample and is then detected by ultrasound transducers, as can be seen in Figure 1.2. These recorded signals can be reconstructed to create a PA image. PA is both non-invasive and non-ionising, and the only major expense is the light source. However, quantifying PA images is a problem that many have struggled with. Several solutions have been proposed, each with their own assumptions and applicability. However, none have achieved results that are both sufficiently accurate and reproducible in a clinical setting.

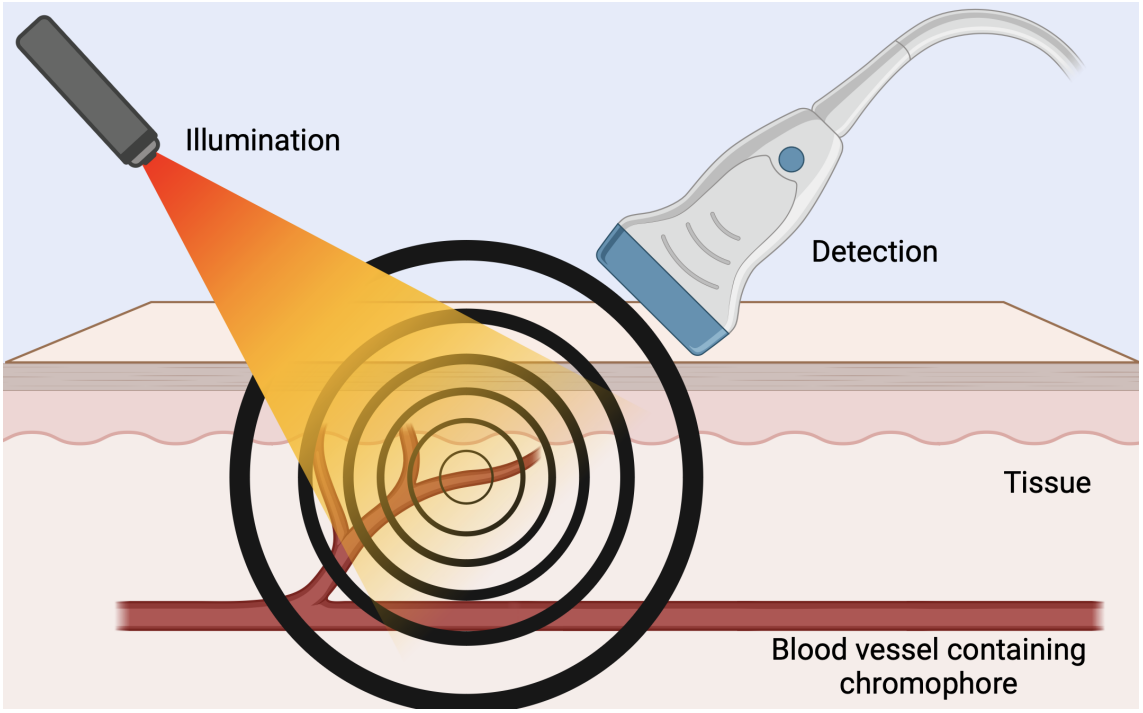

FIGURE 1.2: Schematic representation of the photoacoustic imaging mechanism, created using BioRender.com

1.3 Problem statement

Create an algorithm that is able to utilise a fluence marker during the reconstruction and quantification of tissue composition based on photoacoustic signals.

In this thesis, I will try to answer the following research questions (RQs):

BME RQ: Can the integration of an endogenous reference chromophore, such as blood in the carotid artery, enhance the accuracy of quantitative photoacoustic imaging in the retrieval of non-reference chromophore concentrations, such as in plaques in the carotid artery?

CS RQ: Can the incorporation of virtual detectors (requiring certain optical properties and spatial location as prior) enhance the accuracy of model-based optical inversion in quantitative photoacoustic imaging?

Chapter 2

Theoretical background

Simulating photoacoustic imaging is a complex endeavour. The imaging modality consists of both an optical and an acoustic part, so knowledge of simulating both is a must for accurate results. This chapter is split into these two categories (with a small section in the middle) since the processes do not interact with each other directly.

2.1 Optical simulation

Light exhibits behaviour of both a classical wave and a particle. However, for highly scattering media (such as human tissue [6]) and large spatial scales, modelling light using the assumption of wave-based behaviour will quickly reach the limits of computational simulation. Therefore, modelling based on particle-based behaviour is more common. These models use the radiative transfer equation, also called Boltzmann's transport equation for low-energy, monochromatic photons:

Change in radiance over time — Light source — Change in radiance over space — Absorption and scattering away from direction $\hat{\mathbf{s}}$

$$\frac{1}{c}\frac{\partial \phi(\mathbf{x},\hat{\mathbf{s}},t)}{\partial t} = q(\mathbf{x},\hat{\mathbf{s}},t) - \hat{\mathbf{s}}\cdot\nabla\phi(\mathbf{x},\hat{\mathbf{s}},t) - (\mu_a(\mathbf{x})+\mu_s(\mathbf{x}))\,\phi(\mathbf{x},\hat{\mathbf{s}},t) + \mu_s\int \Theta\left(\hat{\mathbf{s}},\hat{\mathbf{s}}'\right)\phi\left(\mathbf{x},\hat{\mathbf{s}}',t\right)\mathrm{d}\hat{\mathbf{s}}' \tag{2.1}$$

Scattering into direction $\hat{\mathbf{s}}$ from other directions

Here, c is the speed of light in the medium, $\phi(\mathbf{x},\hat{\mathbf{s}},t)$ is the light radiance at position $\mathbf{x}$ in direction $\hat{\mathbf{s}}$ at time t, $q(\mathbf{x},\hat{\mathbf{s}},t)$ is a source of photons, μ_a is the optical absorption coefficient, μ_s is the optical scattering coefficient, and $\Theta\left(\hat{\mathbf{s}},\hat{\mathbf{s}}'\right)$ is the scattering phase function; the probability that a photon that is travelling in direction $\hat{\mathbf{s}}$ will scatter to direction $\hat{\mathbf{s}}'$. This equation does not take wave effects, polarisation, radiative processes, ionisation, inelastic scattering and reactions into account.

Equation (2.1) is generally not solved directly. Instead, its solution can be approximated with several methods. Monte Carlo, often considered to be the golden standard [7], is able to do this with little assumptions while being computationally efficient (can be parallelised, and accelerated using a GPU). Although there are many different implementations, the general concept behind Monte Carlo in optical simulation is that many separate photon packages are propagated throughout a turbid medium, each path randomly determined based on given probability functions [6]. Its error $\mathcal{E}$ depends on the

number of simulated photons N: $\mathcal{E} = O(1/\sqrt{N})$. Since Monte Carlo is a purely numerical technique, inversion and differentiation are complex. Attempts have been made to recover optical properties in tissue using Monte Carlo, but these either resort to iterative forward fitting (such as in Palmer *et al.* [8]) or numerical differentiation (such as in Buchmann *et al.* [9]) which are due to their very nature subject to the noise introduced by separating light propagation into many stochastic light packages.

Since one can assume that acoustic propagation occurs several orders of magnitude slower than heat deposition in pulsed PA imaging (see Section 2.2), it is also possible to assume steady state ($\frac{\partial \phi}{\partial t} = 0$) in Equation (2.1), to find the time-integrated absorbed power density. This results in the time-independent Radiative Transfer Equation (RTE):

$$\left(\hat{\mathbf{s}} \cdot \nabla + \mu_t\right) \phi(\mathbf{x}, \hat{\mathbf{s}}) - \mu_s \int \Theta\left(\hat{\mathbf{s}}, \hat{\mathbf{s}}'\right) \phi\left(\mathbf{x}, \hat{\mathbf{s}}'\right) \mathrm{d}\hat{\mathbf{s}}' = q(\mathbf{x}, \hat{\mathbf{s}}), \tag{2.2}$$

where $\phi(\mathbf{x}, \hat{\mathbf{s}})$ is the time-integrated light radiance, and $\mu_t = \mu_a + \mu_s$ is the total attenuation coefficient. One can integrate the light radiance over all angles $\hat{\mathbf{s}}'$ to obtain the fluence Φ:

$$\Phi(\mathbf{x}) = \int \phi\left(\mathbf{x}, \hat{\mathbf{s}}'\right) \mathrm{d}\hat{\mathbf{s}}', \tag{2.3}$$

which determines the photoacoustic signal (see next chapter). Equation (2.2) is generally speaking not solved directly. One can assume to have no scattering events, which will reduce the equation to $\left(\hat{\mathbf{s}} \cdot \nabla + \mu_a\right) \phi(\mathbf{x}, \hat{\mathbf{s}}) = q(\mathbf{x}, \hat{\mathbf{s}})$. When the illumination (Φ_0) is collimated and planar, and μ_a is constant throughout the tissue, one gets the Lambert-Beer law:

$$\Phi(z) = \Phi_0 \exp(-\mu_a z), \tag{2.4}$$

which describes the exponential decrease of the fluence Φ along the z-direction. However, these assumptions do not hold for the majority of all setups and samples. Instead, it is possible to write the radiance as a sum of spherical harmonics and truncate after N terms. It is often sufficient to take N = 1 when the directional dependence of the light distribution is weak, which is the case for highly scattering tissue: $\mu_s' \gg \mu_a$. The weak directional dependence of the light distribution means that the scattered fluence is almost isotropic, such that one ends up with the Diffusion Approximation (DA):

$$\underbrace{\mu_a \Phi}_{\text{Loss due to absorption}} - \underbrace{\nabla \cdot (D \nabla) \Phi}_{\text{Loss due to diffusion}} = \underbrace{q_0}_{\text{Source term}}, \tag{2.5}$$

where $D = \frac{1}{3(\mu_a + \mu_s')}$ is the diffusion coefficient, $\mu_s' = (1 - g)\mu_s$ is the reduced scattering coefficient, g is the anisotropy parameter, and q_0 is the isotropic directional source term. The DA can be solved numerically, e.g. using finite element methods.

2.2 Photoacoustic effect

With the models of the previous section in mind, it is possible to calculate the light fluence Φ for a given geometry with known optical properties ($\mu_a(\mathbf{x}, \lambda)$ and $\mu_s'(\mathbf{x}, \lambda)$). As mentioned previously, as a result of the presence of naturally occurring chromophores, photons are absorbed. The energy of these absorbed photons will raise the local temperature and thus

the local pressure (initial pressure, denoted p_0). This process is called the photoacoustic effect:

$$H = \Phi \mu_a \tag{2.6}$$

$$p_0 = \hat{\Gamma} H = \hat{\Gamma} \Phi \mu_a. \tag{2.7}$$

Here, H is the absorbed energy, and $\hat{\Gamma}$ is the efficiency of the PA effect. This efficiency depends on the Grüneisen parameter Γ for an absorbing fluid. The resulting initial pressure, in turn, becomes a propagating acoustic wave, which can be detected by ultrasound (US) transducers placed at the boundary of the geometry under study.

There are two main confinement requirements for photoacoustic imaging [10], the thermal and the stress confinement requirements. The former demands that the optical pulse duration τ_p is shorter than the thermal relaxation time $\tau_h = d_c^2/\alpha$:

$$\tau_p < \tau_h, \tag{2.8}$$

where d_c is the dimension of the structure under study and α its thermal diffusivity. The second confinement requirement demands that the optical pulse duration τ_p is shorter than the stress relaxation time $\tau_s = d_c/v_s$:

$$\tau_p < \tau_s, \tag{2.9}$$

where v_s is the velocity of sound. The better these requirements are fulfilled, the higher the PA efficiency (with an upper limit of Γ).

To quantify this PA efficiency $\hat{\Gamma}$, one must look at several factors. Part of the absorbed energy in PA imaging is used to change the mass density, and the other part is transformed in an increase in heat (and thus pressure). The ratio between these determines the PA efficiency, the more energy that goes to the latter, the higher the efficiency. First, in a fluid with homogenous thermodynamic properties, the rate of change in the local temperature T', the pressure p' and the density ρ' can be expressed as [5]:

$$\frac{\partial p'}{\partial t} = \left(\frac{1}{\rho k_T}\right) \frac{\partial \rho'}{\partial t} + \left(\frac{\beta}{k_T}\right) \frac{\partial T'}{\partial t}. \tag{2.10}$$

Here, ρ is the ambient mass density, k_T is the isothermal compressibility, and β is the volume thermal expansivity of the fluid. Given that the PA process takes place where the temperature confinement requirement is fulfilled, the first term can be neglected. Furthermore, given that the absorbed power density $\mathcal{H}(\mathbf{x}, t)$ is the product of the absorbed energy $H(\mathbf{x})$ and the temporal shape function $f(t)$: $\mathcal{H}(\mathbf{x}, t) = H(\mathbf{x}) f(t)$, the temperature increase rate can be written as: $\partial T'/\partial t = \mathcal{H}/\rho C_v$ where C_v is the specific heat capacity at constant volume [5]. Therefore, the increase in local pressure can be related to the absorbed optical power density:

$$\frac{\partial p'}{\partial t} = \left(\frac{\beta}{\rho C_\nu k_T}\right) \mathcal{H}(\mathbf{x}, t) = \Gamma H(\mathbf{x}) f(t), \tag{2.11}$$

where the Grüneisen parameter is denoted as $\Gamma = \beta/(\rho C_v k_T)$ which is a thermodynamic material property. If one integrates this equation over the duration of the optical pulse, one gets $p_0 = \Gamma H$, which is a special case of Equation (2.7) where $\hat{\Gamma} = \Gamma$. This is only valid when the absorber and the propagation medium have the same thermodynamic material properties and if the confinement requirements are met, otherwise $\hat{\Gamma} < \Gamma$.

2.3 Acoustic simulation

While light can be seen as both a classical wave and a particle, sound does not exhibit this wave-particle duality. Instead, the propagation of photoacoustic pressure waves can be described by a wave equation. Since the initial pressure in photoacoustics is typically in the order of 10 kPA or less [11], the time evolution of these waves can be modelled using the equations of linear acoustics. Moreover, for soft biological tissue, one can assume that the propagation medium is isotropic and quiescent, that the pressure flow is irrotational and that shear waves can be neglected. Thus, the behaviour of photoacoustic waves can be captured in the first-order equations of motion, continuity, and state for a lossless medium, respectively:

$$\text{momentum conservation:} \quad \frac{\partial \mathbf{u}}{\partial t} = -\frac{1}{\rho_0}\nabla p, \tag{2.12a}$$

$$\text{mass conservation:} \quad \frac{\partial \rho}{\partial t} = -\rho_0 \nabla \cdot \mathbf{u}, \tag{2.12b}$$

$$\text{pressure-density relation:} \quad p = v_{s,0}^2 \rho, \tag{2.12c}$$

where the initial conditions are given by Equation (2.7) and $\frac{\partial p_0}{\partial t} = 0$. Here, $\mathbf{u}$ is the acoustic particle velocity, ρ_0 is the ambient density, ρ is the acoustic density, $v_{s,0}$ is the isentropic sound speed, and p is the acoustic pressure. These equations can be combined to get a second-order photoacoustic wave equation [12]:

$$\left(\nabla^2 - \frac{1}{v_s^2}\frac{\partial^2}{\partial t^2}\right) p(\mathbf{r}, t) = -p_0(\mathbf{r})\frac{d\delta(t)}{dt}, \tag{2.13}$$

where the acoustic wave $p(\mathbf{r}, t)$ at position $\mathbf{r}$ at time t is the result of the initial pressure $p_0(\mathbf{r})$, travelling at a velocity v_s as a result of a $\delta(t)$ optical pulse.

The first-order partial differential equations denoted in Equations (2.12) can be solved through the usage of many numerical methods. Applying classical finite element and finite difference approaches generally requires a minimum of 10 grid points per acoustic wavelength to achieve an acceptable level of accuracy here, which can quickly result in hundreds of GB of computer memory for this application [13]. Therefore, to minimise the number of time steps and memory required, other methods are necessary. One such approach uses the k-space pseudospectral method. This approach combines a temporal propagator expressed in the spatial frequency domain or k-space, with the spectral calculation of spatial derivatives. The latter can be calculated using the Fourier collocation method. This approach is discussed briefly, as an extensive explanation is beyond the scope of this thesis.

Spatial gradients can be determined numerically through the usage of standard finite-difference schemes, where they are determined locally based on function values at neighbouring grid points. This starts with linear interpolation between two function values but can include more neighbouring points by fitting a higher-order polynomial and taking its derivative. In general, using more points gives a more accurate determination. Using spectral calculations to determine spatial derivatives, one does not fit a polynomial to known neighbouring points but fits a Fourier series to all known data [14]. The major advantages of using this Fourier collocation spectral method are that A) it is possible to utilise the speed that comes with the fast Fourier transform method and B) the basis functions are sinusoidal, so neither six nor ten grid points per acoustic wavelength are required, but only two.

The Fourier collocation spectral method is only able to improve efficiency in the spatial domain, in the temporal domain other methods are required. However, conventional finite

difference schemes introduce errors into the solution which can only be controlled through the limitation of the time-step size, which is undesirable. To overcome this limitation, k-space methods can be applied. This means taking the spatial Fourier transform of the known equation and discretising the resulting time derivative using a second-order central difference. The difference between the standard finite difference and the exact approximation is reduced to an operator in the spatial frequency domain, also called the k-space operator [15]:

$$\kappa = \operatorname{sinc}\left(v_{s,\mathrm{ref}} k \Delta t/2\right), \tag{2.14}$$

where $v_{s,\mathrm{ref}}$ is a scalar reference sound speed.

Using this approach, it is possible to rewrite Equation (2.12a) and (2.12b) as:

$$\frac{\partial}{\partial \xi} p^n = \mathcal{F}^{-1}\left\{ik_\xi \kappa e^{ik_\xi \Delta\xi/2} \mathcal{F}\left\{p^n\right\}\right\}, \tag{2.15a}$$

$$u_\xi^{n+\frac{1}{2}} = u_\xi^{n-\frac{1}{2}} - \frac{\Delta t}{\rho_0} \frac{\partial}{\partial \xi} p^n, \tag{2.15b}$$

$$\frac{\partial}{\partial \xi} u_\xi^{n+\frac{1}{2}} = \mathcal{F}^{-1}\left\{ik_\xi \kappa e^{-ik_\xi \Delta\xi/2} \mathcal{F}\left\{u_\xi^{n+\frac{1}{2}}\right\}\right\}, \tag{2.15c}$$

$$\rho_\xi^{n+1} = \rho_\xi^n - \Delta t \rho_0 \frac{\partial}{\partial \xi} u_\xi^{n+\frac{1}{2}}. \tag{2.15d}$$

Equations (2.15a) and (2.15c) are the result of spatial gradient calculations using the Fourier collocation spectral method. Equations (2.15b) and (2.15d) are update steps, based on a k-space correct first-order accurate forward difference. In these equations, $\xi = x, y$ for the 2-dimensional case, $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the forward and inverse spatial Fourier transform respectively, i the imaginary unit, k_ξ the wavenumber and $\Delta\xi$ the grid spacing in the ξ direction, Δt is the time step, and n and $n+1$ the current and next time step. The attentive reader might notice $n \pm \frac{1}{2}$, which denote time-staggered points, arising because the update steps are interleaved with the gradient calculations. With these equations in mind, it is finally possible to calculate the pressure for the next time step:

$$\rho^{n+1} = \sum_\xi \rho_\xi^{n+1}, \tag{2.16}$$

$$p^{n+1} = v_{s,0}^2 \rho^{n+1}. \tag{2.17}$$

By repeatedly applying these equations, one can simulate the propagation of pressure throughout the tissue.

Chapter 3

Related Work

In PA, tissue is illuminated, resulting in a pressure increase, which in turn can be detected at the boundary of the sample. This process is shown in Figure 3.1A. To determine the tissue composition of the sample under study using the measured acoustic time series, one needs to invert these steps, see Figure 3.1B. Starting with the measured time series $p(t)$, the acoustic operator $A(\cdot)$ is inverted so that the reconstructed pressure is acquired: $p_{recon} = A^{-1}(p(t))$. This is called the acoustic inverse problem. Next, the effect of the PA efficiency $\hat{\Gamma}$ is removed so that the absorbed energy distribution is acquired: $H = \hat{\Gamma}^{-1}(p_{recon})$. After this, the optical operator $T(\cdot)$ is inverted, such that the optical properties of the tissue are acquired: $(\mu_a, \mu_s')) = T^{-1}(H)$. This is called the optical inverse problem. If all of these steps are repeated for separate illumination with several optical wavelengths, the concentration of each chromophore can be calculated. Together, this forms the composition inverse problem. The acoustic, optical, and composition inverse problems are now discussed in turn.

3.1 Acoustic inverse problem

Several acoustic reconstruction techniques are commonly used. These techniques are first examined, and then their limitations and how to remedy them are discussed.

3.1.1 Reconstruction techniques

In this section the following reconstruction techniques are discussed: universal backprojection, delay-and-sum, and time reversal.

Universal backprojection, one of the more elegant forms of acoustic reconstruction, offers exact reconstruction equations for specific imaging geometries, based on Equation (2.13) and several assumptions [12]. E.g. Xu *et al.* developed a universal backprojection equation for the planar, cylindrical and spherical imaging geometry [12]. These equations can be extended to include the limited-angle view case or the (somewhat) acoustic inhomogeneity of the sample under study. While computational intensive, they have proven to have their practical applications [16]. The backprojection scheme can be modified to include filtering, where either the measured acoustic time series is filtered before backprojection, or the reconstructed pressure is filtered after backprojection. Often, a low-pass filter is chosen to dampen high-frequency noise [17].

Another typical reconstruction technique is delay-and-sum (DAS). The main idea behind DAS is to shift the recorded time signals with appropriate time *delay*, and *sum* them

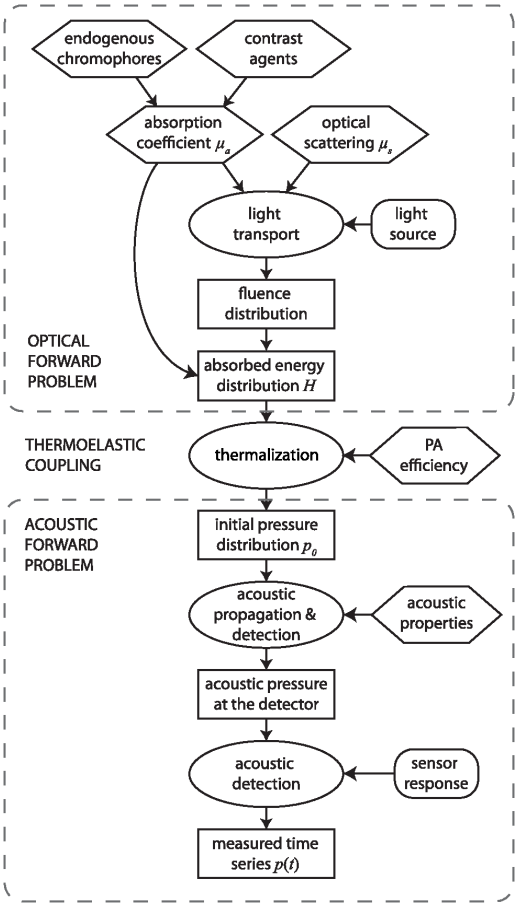

(A) PA signal generation: starting with chromophore concentrations, ending with the measured acoustic time series $p(t)$

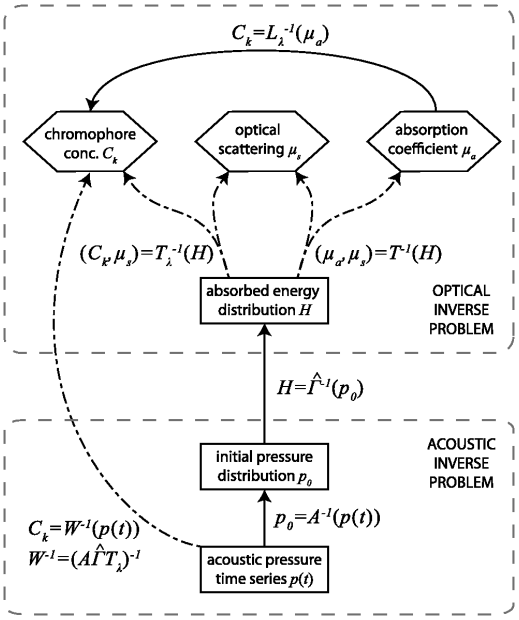

(B) The inverse of Figure 3.1A: reconstruction of chromophore concentrations, starting with measured acoustic time series $p(t)$

FIGURE 3.1: Forward and inverse problems in quantitative PA imaging. Both reproduced from Cox *et al.* [5]

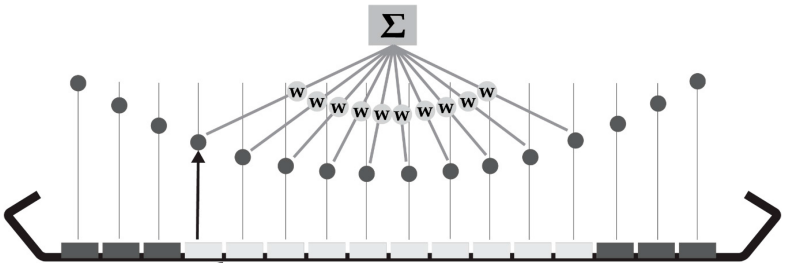

FIGURE 3.2: Schematic overview of delay-and-sum (DAS) beam-forming, from Hendriks *et al.* [20]. If required, the signals can be weighted w by an apodization function

together [18], see Figure 3.2. In equation this is denoted by:

$$p_{DAS}(t) = \sum_{i=1}^{M} p_i \left(t - \tau_i\right), \tag{3.1}$$

where the detected time signals $p_i(t)$ from the M array elements are delayed by a time delay of τ_i and are aggregated to get the output $p_{DAS}(t)$ [19]. However, this approach is prone to high levels of sidelobing and a low resolution, since it suffers from weak suppression of interfering signals and low levels of off-axis signal rejection. To mitigate these problems, one can use delay-multiply-and-sum (DMAS), where the delayed measurement points are multiplied before the summation:

$$p_{DMAS}(t) = \sum_{i=1}^{M-1} \sum_{j=i+1}^{M} p_i \left(t - \tau_i\right) p_j \left(t - \tau_j\right). \tag{3.2}$$

Even more complex variations are possible, such as the double-stage delay multiply and sum [18].

A different reconstruction technique for acoustic inversion is time reversal. As the name suggests, the initial pressure is reconstructed using the measured acoustic time series (from $t = 0$ to $t = T$) in reverse order (from $t = T$ to $t = 0$), as a time-varying Dirichlet boundary condition at the respective detector locations [21]. The waves introduced by the imposed boundary conditions are propagated into the sample under study (which has zero initial conditions) by a forward propagation model such as those introduced in Section 2.3. After time T the propagation is stopped, and the current acoustic pressure within the sample then forms the reconstruction. The choice of model and model settings (such as the size of the time step) dictate the accuracy of the reconstruction. Since this scheme is in essence solving an inverse problem by solving a forward problem, it is relatively straightforward to include sample heterogeneities and define complex imaging geometries.

3.1.2 Limitations

Although each reconstruction technique has its own inherent limitations, a perfect reconstruction technique will still need to correct for the limitations of the US transducer. Most transducers used in PA suffer from the following issues: a limited view, a bandlimited nature, directivity of its elements, and of course noise. These issues can and will lead to a reduction in accuracy in the reconstruction and should therefore be mitigated. Several approaches will be discussed.

Iterative reconstruction

A method to compensate for limitations such as limited view is iterative reconstruction. This method can be applied with any acoustic reconstruction technique that is able to solve both the forward $A(\cdot)$ and the inverse $A^{-1}(\cdot)$ acoustic problem. For a given ground truth initial pressure p_0, its measured acoustic time series $p(t)$ can be reconstructed such that the reconstructed initial pressure is $p_{recon}^{(k=0)} = A^{-1}(p(t))$. By iteratively 1) solving the forward problem, 2) subtracting the resulting time series from the measured time series $p(t)$, 3) solving the inverse problem for the resulting difference, and 4) adding that to the previous estimate of the reconstructed pressure, one can improve the accuracy of the reconstruction [22]. This can be denoted as:

$$p_{recon}^{(k+1)} = p_{recon}^{(k)} + A^{-1}\left(p(t) - A(p_{recon}^{(k)})\right) \tag{3.3}$$

This approach can be extended to include a positivity condition: after each iteration replace all negative values with zero, since the initial pressure can never be negative. Another extension is the usage of a scaling factor $\eta < 1$ with which to multiply the addition step, which can be fixed or variable during the process.

Deconvolution

Transducers are not ideal, and their finite bandlimited nature distorts the recorded signal. As described by Wang *et al* [23], the recorded pressure $p_r(t)$ is the result of a convolution of the arriving pressure $p_a(t)$ with the transducer impulse response $h(t)$:

$$p_r(t) = p_a(t) * h(t), \tag{3.4}$$

when a point detector is assumed. To optimise reconstructions, one must minimise the effect of this distortion on the recorded signal.

A most obvious approach would be to remove this effect using Fourier transformations [24]:

$$p_a(t) = \mathcal{F}^{-1}\left[\frac{\mathcal{F}\left[p_r(t)\right] W(\omega)}{\mathcal{F}\left[h(t)\right]}\right], \tag{3.5}$$

Here, $\mathcal{F}[\cdot]$ and $\mathcal{F}^{-1}[\cdot]$ denote the forward and inverse temporal Fourier transformation. One uses the window function $W(\omega)$ to prevent unwanted frequencies from being amplified. This could look like $W(\omega) = |\sin(\pi\bar{\omega}/2)|$ where $\bar{\omega} = \frac{\omega}{\omega_{\max}} \in [-1, 1]$. This approach can be extended, such as Wiener deconvolution, which is more resistant to noise [24].

An alternative is the matrix deconvolution method [24], paired with a regularisation technique. One can assume that each recorded signal consists of a combination of impulse responses, each with a different intensity and delay. Therefore, it is possible to rewrite Equation (3.4) as a matrix equation:

$$p_r = C\ p_a, \tag{3.6}$$

where p_r and p_a are both column vectors containing $p_r(t)$ and $p_a(t)$ for each recorded time respectively, and C is a non-circulant matrix. C contains increasingly shifted copies of the impulse response $h(t)$ along its columns, such that the origin of two neighbouring columns differs with $dx = v_s\ dt$ where v_s is the speed of sound and dt the time between two

neighbouring measurements from a transducer element. Since C is expected to be rank-deficient, one can apply L-2 norm regularisation. This means minimising the following cost function:

$$\min_{p_a} ||C\ p_a - p_r||^2 + ||\Lambda\ p_a||^2, \tag{3.7}$$

where the matrix Λ determines the nature of the regularisation, e.g. $\Lambda = \beta I$, where β is a scalar setting the regularisation strength, and I the identity matrix. The regularised estimate $\hat{p}_a$ can be found by solving Equation (3.7):

$$\hat{p}_a = \left(C^T C + \Lambda^T \Lambda\right)^{-1}\ C^T\ p_r. \tag{3.8}$$

Another regularisation technique that can be used to normalise is L-1 norm, which means minimising the following cost function:

$$\min_{p_a} ||C\ p_a - p_r||^2 + ||\Lambda\ p_a||. \tag{3.9}$$

This equation does not have an analytical solution, so algorithms such as Coordinate Descent [25] or an algorithm based on Alternating Direction Method [26] of Multipliers are required.

Deconvolution extended

Up to this moment, $p_r(t)$ of each transducer element had to be deconvoluted separately before reconstructing the initial pressure. However, deconvolutions and reconstruction can be combined into one deconvolution [27]. To this end, an imaging region of $n \times n$ points is defined. Next, this region is vectorised into a column vector x with length n^2 by stacking the columns. Following this, the system's response to each entry in the vectorised image is recorded. For each entry, the resulting $p_r(t)$ (one per transducer element) are stacked into one long column vector with length m, and these vectors are combined into one matrix A with dimensions $m \times n^2$. From this, one can determine p_0 directly by solving $Ax = b$, where b is the vectorisation of a measurement with the system. Solving this equation is likely ill-posed, making regularisation (such as L-1 and L-2 as mentioned before) suitable. However, this approach is very computationally intensive and requires a close match between simulation and measurement.

3.1.3 Deep learning

Of course, all the approaches mentioned so far are all model-based rather than (partly) data-driven. In this section, deep learning methods that can be used in the acoustic inverse problem will be discussed. A large part of research in applying deep learning to the acoustic inverse problem has been focused on A) enhancing existing model-based methods, or B) performing the entire acoustic inverse problem at once.

Enhancing existing model-based methods can, for example, be done by training weights in a universal backprojection algorithm [28], training an additional regulariser in an iterative reconstruction scheme [29], or integrating a Convolutional Neural Network (CNN) for each iteration step in the reconstruction scheme [30].

On the other hand, deep-learning methods can also be deployed to solve the entire acoustic inverse problem. CNNs have been used [31], as well as a modified U-Net architecture [32], in combination with a lookup table [33], and generative adversarial networks [34]. Although deep learning methods can offer an increase in speed in the order of magnitude,

they often lack acceptable levels of robustness, uncertainty awareness, and generalisation to experimental data [35].

Finally, deep learning can also be applied for artefact removal and overall image quality enhancement, where well-established methods from other imaging fields can be of direct benefit [35].

3.2 Optical inverse problem

Assuming an accurate solution to the acoustic inverse problem and a known PA efficiency, which results in knowing H accurately, solutions to the optical inverse problem are now required.

A simple case would be a small perturbation in μ_a over a known and homogenous $\mu_{a,0}$ and $\mu'_{s,0}$ [5]: $\mu_a(\mathbf{x}) = \mu_{a,0} + \delta\mu_a(\mathbf{x})$. The fluence that is calculated without this perturbation can be denoted as $\Phi_0(\mathbf{x})$. The actual fluence then is a combination of $\Phi_0(\mathbf{x})$ and $\delta\Phi(\mathbf{x})$. The difference in absorbed energy (while neglecting the second-order term) is then:

$$\delta H(\mathbf{x}) \approx \delta\mu_a(\mathbf{x})\Phi_0(\mathbf{x}) + \mu_{a,0}\delta\Phi(\mathbf{x}). \tag{3.10}$$

Here, the first term is the result of a difference in the local absorption coefficient, and the second term is the result of a change in the local fluence, caused by the absorption perturbation. Still, the change in absorbed energy $\delta H(\mathbf{x})$ is in reality non-linearly related to the perturbation in the absorption coefficient ($\delta\Phi$ depends on $\delta\mu_a$). One can assume that $\delta\Phi = 0$, which will reduce Equation (3.10) to:

$$\delta H(\mathbf{x}) \approx \delta\mu_a(\mathbf{x})\Phi_0(\mathbf{x}). \tag{3.11}$$

By determining $\delta H(\mathbf{x})$ and $\Phi_0(\mathbf{x})$, one is now able to determine $\delta\mu_a(\mathbf{x})$. However, this assumption is almost exclusively suitable for very weak absorbers.

In most cases, the absorption cannot be modelled as a small perturbation on a homogeneous background. Instead, one can reformulate Equation (2.6) as a fixed-point iteration without doing linearising assumptions [36]:

$$\mu_a^{(n+1)}(\mathbf{x}) = \frac{\hat{H}(\mathbf{x})}{\Phi^{(n)}\left(\mathbf{x}, \mu_a^{(n)}(\mathbf{x})\right)}, \tag{3.12}$$

where the absorption coefficient estimate $\mu_a^{(n)}(\mathbf{x})$ at iteration n results in the fluence estimate $\Phi^{(n)}$, and $\hat{H}(\mathbf{x})$ is the measured absorbed energy distribution. Of course, measuring $\hat{H}(\mathbf{x})$ can only be done by accurate acoustic reconstruction and accurate knowledge of the PA efficiency. Both parts can introduce noise ($\mathcal{N}$), so the measured absorbed energy distribution $\hat{H}$ will look more like $\hat{H}(\mathbf{x}) = H(\mathbf{x}) + \mathcal{N}$. This will introduce instability at places where Φ is small, to which a regularisation parameter σ can form a solution:

$$\mu_a^{(n+1)}(\mathbf{x}) = \frac{\hat{H}(\mathbf{x})}{\Phi^{(n)}\left(\mathbf{x}, \mu_a^{(n)}(\mathbf{x})\right) + \sigma(\mathbf{x})}. \tag{3.13}$$

This regularisation parameter can be a constant, or dependent on $\mathbf{x}$, the iteration number n, and/or the signal-to-noise ratio (SNR).

Another approach starts by looking at the DA (Equation (2.5)) again. One can see that $\mu_a\Phi$ is directly contained in the DA. Therefore, it can be rewritten as [37]:

$$\nabla \cdot (D\nabla)\Phi = \mu_a\Phi - q_0 = H - q_0. \tag{3.14}$$

Assuming that $D \approx \frac{1}{3\mu_s'}$, the equation can be solved numerically and the absorption coefficients can be extracted. If one does not assume that D is independent of μ_a, an iterative approach can be used to estimate μ_a.

Up until now, all approaches shown for the optical inverse problem have been formulated as derivative-free optimisation problems. However, gradient-based methods are superior in many cases due to their (possibility of) speed and accuracy. To formulate such a method [38], an error function is determined:

$$\mathcal{E} = \frac{1}{2} \int \left(\hat{H} - H\left(\mu_a, \mu_s'\right) \right)^2 dV. \tag{3.15}$$

Differentiating this equation with respect to μ_a and μ_s' at a single point gives:

$$\frac{\partial \mathcal{E}}{\partial \mu_a} = -\Phi \left(\hat{H} - H \right) + \Phi\Phi^*, \text{and} \tag{3.16}$$

$$\frac{\partial \mathcal{E}}{\partial \mu_s'} = -3\kappa^2 \nabla\Phi^* \cdot \nabla\Phi, \tag{3.17}$$

where

$$\left(\mu_a + \nabla \cdot \kappa\nabla\right)\Phi^* = \mu_a \left(\hat{H} - H \right). \tag{3.18}$$

Adjoint Equation (3.18) is comparable to the DA (Equation (2.5)), the only difference is that this equation has a different source term, making it solvable with the same methods as used to solve the DA. By calculating these gradients, it is possible to find the solution (a μ_a and a μ_s' distribution) that minimises the error function.

3.2.1 Deep learning

Just like in the acoustic inverse problem, deep learning methods can offer an improvement over model-based methods in the optical inversion problem. Several approaches have been proposed, all using the U-Net architecture. Cai *et al.* showed that ResU-net, a novel end-to-end deep neural network, was able to determine optical properties in certain settings [39]. A different paper has also shown that U-Net models can estimate fluence and optical properties in certain settings, but also developed a method to estimate the uncertainty in the resulting output [40]. Furthermore, a different U-Net model was trained and tested on synthetic and experimental data from simple phantoms [41]. However, only this small amount of research has been published in this region, none of which has been verified *in vivo*, most likely due to being one of the most challenging problems in quantitative PA imaging [35].

3.3 Composition inverse problem

With the solutions to both the optical and acoustic inverse problem in mind, the composition inverse problem is now discussed: determination of the concentrations of chromophores based on measured optical properties.

A common assumption is that the optical properties at a given optical wavelength λ are a linear combination of the optical properties of the chromophores present at a given point $\mathbf{x}$:

$$\mu_a(\mathbf{x}, \lambda) = \sum_{k=1}^{K} C_k(\mathbf{x}) a_k(\lambda), \tag{3.19}$$

or equally:

$$\begin{pmatrix} \mu_a(\mathbf{x}, \lambda_1) \\ \vdots \\ \mu_a(\mathbf{x}, \lambda_N) \end{pmatrix} = \begin{pmatrix} \alpha_1(\lambda_1) & \dots & \alpha_K(\lambda_1) \\ \vdots & \ddots & \vdots \\ \alpha_1(\lambda_N) & \dots & \alpha_K(\lambda_N) \end{pmatrix} \begin{pmatrix} C_1(\mathbf{x}) \\ \vdots \\ C_K(\mathbf{x}) \end{pmatrix}, \tag{3.20}$$

for K chromophores, where a_k is the molar absorption coefficient and C_k the concentration of chromophore k. If one knows the molar absorption coefficient spectra of all chromophores present, the inverse problem is to solve the system of linear equations in Equation (3.20): solve $\mathbf{b} = A\mathbf{x}$ for x. This matrix equation can either be solved directly or approximated, for example, with `mldivide()` in MATLAB, $\mathbf{x} = A^{-1}\mathbf{b}$. Therefore, recovering the composition means determining and unmixing μ_a linearly for multiple optical wavelengths, comparable to determining blood oxygen saturation with pulse oximetry [42].

Chapter 4

Quantitative Photoacoustic Imaging

In this chapter, a novel approach to quantify photoacoustic imaging using a fluence marker will be introduced, detailed, and quantified, which henceforth will be called the virtual detector method. Ideally, this method would have the measured acoustic time series as an input. However, the method will have simulated absorbed energy maps as input, so the improvement in optical inversion can be demonstrated and proof of principle shown. This means it is assumed that the acoustic inversion is done accurately and with negligible distortion. Future work can relatively easily expand the method as to include acoustic inversion, seeing as the acoustic methods are already included in the code, which is available upon request. For demonstration purposes, in Figure 4.1 a simulation of the optical forward, acoustic forward and backward problem on a simple phantom is included. The optical simulation here is done using the forward engine of NIRFAST (with `femdata_FD()` and `BiCGStab_GPU`) [43], which is an implementation of the Diffusion Approximation (see Equation (2.5)). The acoustic simulation is done using k-Wave (with `kspaceFirstOrder2D()`) [21], including the deconvolution scheme from Equation (3.8) and the iterative scheme from Equation (3.3).

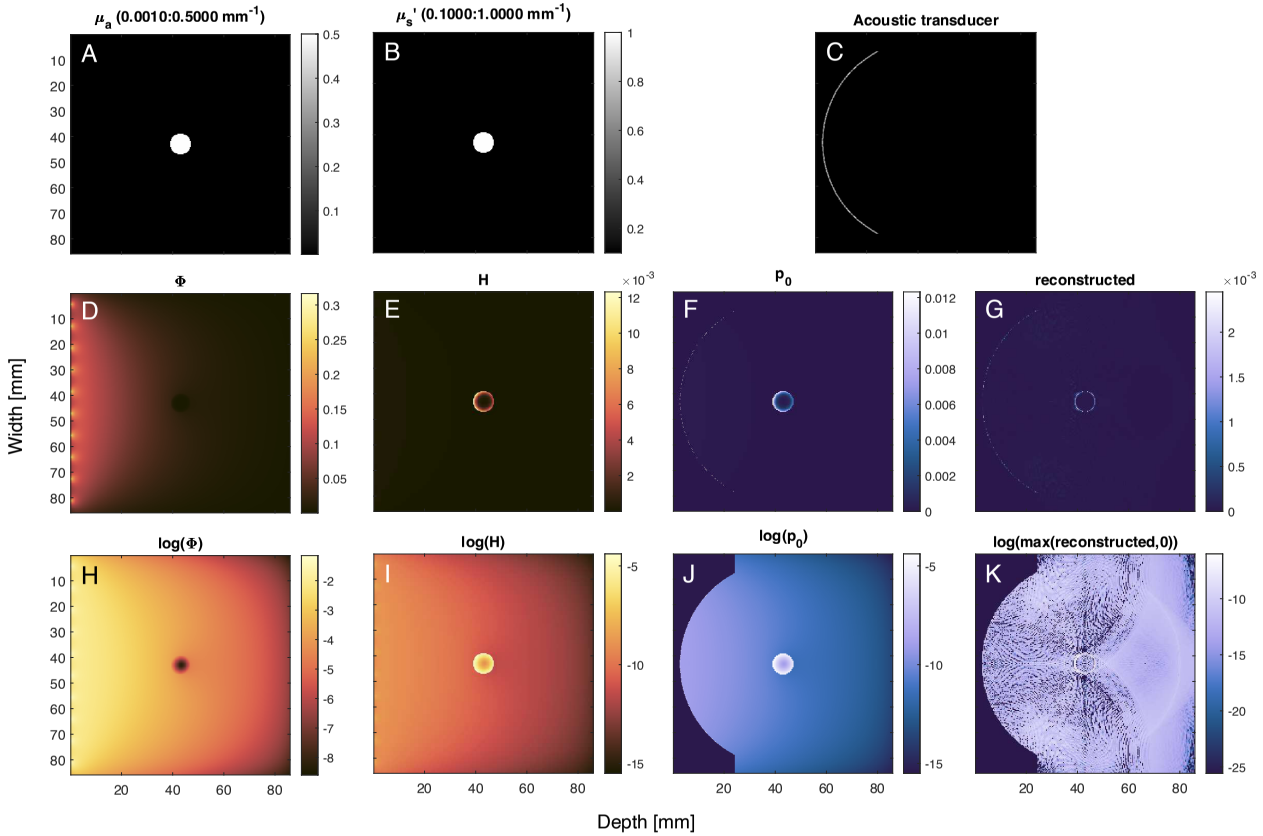

FIGURE 4.1: Demonstration of solving the optical forward, acoustic forward and backward problem. A and B: optical properties of the simple phantom, absorption and reduced scattering coefficient distribution respectively, C: placement of curved acoustic transducer as used in the acoustic forward and backward problem, D and H: calculated fluence using NIRFAST based on optical properties in A and B, E and I: absorbed energy distribution ($H = \Phi\mu_a$), F and J initial pressure ($p_0 = \hat{\Gamma}H$), assuming $\hat{\Gamma} = 1$, G and K: reconstructed pressure, forward and backward acoustic problem solved using k-Wave, backward including the deconvolution scheme from Equation (3.8) and the iterative scheme from Equation (3.3). Properties of the acoustic transducer: number of detector elements = 256, radius of curve = 40 mm, opening angle of curve = 125 degrees, centre frequency = 4.9 MHz, and bandwidth = 80 %.

4.1 Introduction

In the rest of this chapter, a novel method that employs a reference chromophore within the tissue as a fluence marker, to enhance absolute quantification in photoacoustic imaging, is proposed. The hypothesis is that prior knowledge of the optical properties and location of a chromophore within the tissue can aid in refining quantitative photoacoustic imaging techniques. A practical application of this method is the quantification of plaque in the carotid artery, where blood in the lumen can serve as the reference chromophore with known optical properties, or in scenarios involving a known contrast agent injected into the tissue [44].

The methods in quantitative PA imaging are aimed at obtaining chromophore concentrations from PA images at two or more wavelengths. The complex nature of combined optical and acoustic inversion makes it troublesome to quantify and requires both accurate optical and acoustic inversion to obtain the chromophore concentration. The lack of quantitative imaging is a major hindrance in clinical translation [45]. Absolute chromophore quantification has had limited success in oxygen saturation imaging where system-related unknown parameters can be cancelled out [46].

The approach is investigated in the context of existing quantitative PA imaging algorithms. Namely, by comparing the novel method with two algorithms proposed by Ben Cox *et al.* for absorption coefficient retrieval [36, 38]. The goal is to see whether giving the optical properties of one chromophore as a prior can improve the quantification. Given that most photoacoustic systems can also obtain ultrasound images, it is possible to utilise the spatial information of a reference chromophore, such as arterial blood in carotid imaging. This can be compared to the research of Maslov *et al.* [47]. Here, they used a plain black absorbing film with a spectrally flat absorption coefficient to measure the spectrum of the PA signal generated by it to rectify the spectral colouring. Furthermore, Rajian *et al.* used an exogenous contrast agent with known optical properties to measure the local fluence [44]. In this chapter, this approach of utilising a reference chromophore in combination with quantitative PA imaging algorithms by constraining the spatial location and the optical properties of the reference chromophore is extended. This is proposed in the context of carotid artery imaging as to quantify plaque composition.

The approach of reference chromophore as a fluency marker will be demonstrated, and the methods on several phantoms will be tested. The implementation of the novel algorithm is done using the open-source NIRFAST toolbox so that the method can be widely used. Furthermore, the approach for the target application of carotid plaque imaging using simulated carotid phantoms is demonstrated.

4.2 Methods

4.2.1 Quantitative Photoacoustic Imaging

The difficulty in quantifying photoacoustic imaging can mainly be attributed to two separate problems: optical and acoustic inversion. Starting with the measured time series $p(t)$, one needs to reconstruct the original initial pressure p_0 (acoustic inversion). This reconstructed pressure cannot directly be converted into a concentration, instead one needs to correct for the fluence Φ to determine the optical absorption coefficient μ_a (optical inversion), which can be converted into concentration. There are certainly several other things that hamper quantification, such as the PA efficiency and the limited sensor response. However, as the optical inversion is the most challenging and affects the outcome the

most, accurate acoustic inversion is assumed, so the improvement in the optical inversion can be investigated.

For a given point x inside a sample, the fluence $\Phi(x)$ is determined by A) the illumination at the boundary and B) the absorption and scattering of molecules surrounding it. Under certain conditions this light propagation can be described by the Diffusion Approximation (DA), see Equation (2.5). The DA can be solved and inverted numerically, e.g. using finite element methods. The DA and its inversion are used in NIRFAST, a modelling and image reconstruction package developed for Diffuse Optical Tomography (DOT) [48]. DOT, as with all tomography techniques, requires illumination and detection at different angles. DOT and thus per extension NIRFAST require optical sources and optical detectors placed at the boundary of the sample. In photoacoustic imaging, optical sources are present, but optical detectors are not.

4.2.2 Reference chromophore approach

The final goal is to determine the concentration of several chromophores present in the sample. To do this, fluence correction is crucial. Many techniques have been developed over the years, but none are perfect. In this thesis, the idea of a reference chromophore is further developed. Specifically for the target setup, blood as an endogenous reference chromophore can be used since its optical properties are quite well known, especially if oxygen saturation is measured. The absorption coefficient of blood can be accurately determined by measuring the blood oxygen saturation of the patient under study (e.g. using pulse oximetry) in combination with linear interpolation and a lookup table such as the one presented by Bosschaart *et al.* [49]. The fundamental idea here is that, if the absorbed energy is determined accurately, and the optical absorption coefficient of the reference chromophore is known, the fluence can be calculated there, which gives us a fluence marker. This measured fluence in the reference chromophore can be used during the iterative reconstruction of the absorption coefficients, such that the difference between measured and calculated fluence is minimised.

4.2.3 Virtual detector method

This approach starts with measuring the PA signal. This pressure time series $p(t)$ is first reconstructed to get back the initial pressure, which is then converted to the absorbed energy H by dividing the pressure by the known PA efficiency factor $\hat{\Gamma}$. Next, the measured fluence Φ^M is determined at the inner boundary of the blood vessel (as the fluence quickly decays inside volumes with strong absorbing chromophores). In the simulation, a setup with comparable illumination as to the measurement is created. The optical absorption coefficient is initialised using one iteration of the direct inversion (see next section) where the pixels that belong to the reference chromophore are replaced with the known absorption coefficient. The optical reduced scattering is assumed to be known. With this initial guess in place, the fluence distribution is calculated using the DA and the calculated fluence Φ^C at the inner boundary of the blood vessel is extracted. Since the initial guess is presumably wrong, there will be a difference between Φ^M and Φ^C for each blood pixel. Using this error, the μ_a value at each pixel is updated, so the error will decrease. The μ_a values of the reference chromophores are fixed, as these are already correct and should therefore not be changed. Φ^C is calculated for the new μ_a distribution, and the error is determined. This cycle is repeated until the error stops decreasing or a maximum number of iterations is reached. An overview of this method can be seen in Figure 4.2. By determining the μ_a distribution for several wavelengths, the concentration of the chromophores can be

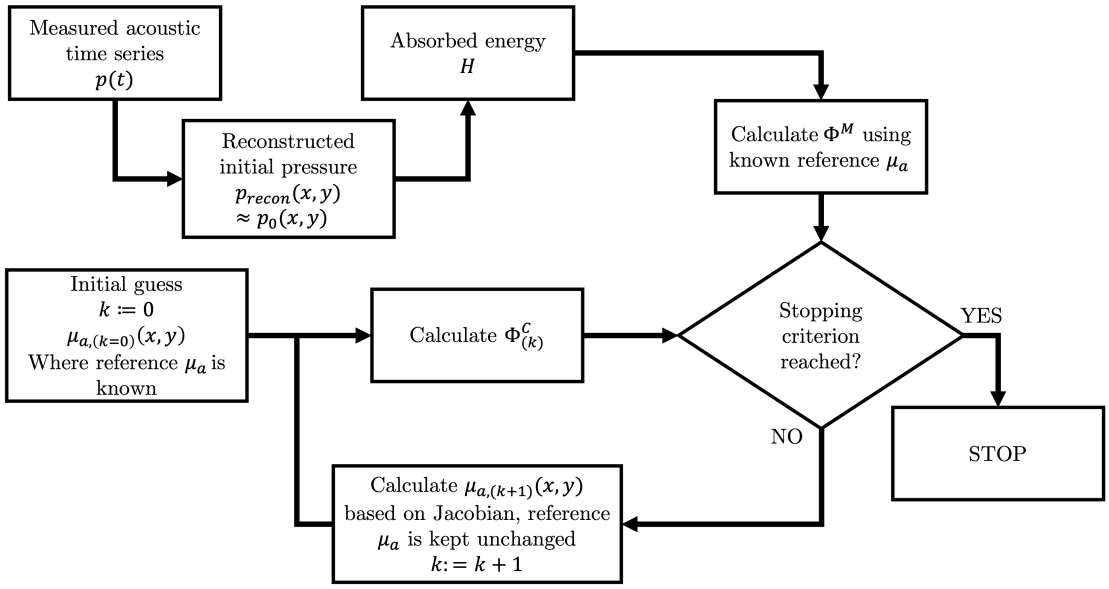

FIGURE 4.2: Overview of the virtual detector method. One starts with solving the acoustic backward problem using the measured acoustic time series and then determines the measured fluence at the virtual detector. Next, one sets an initial guess, calculates the fluence, and examines if the stopping criterion is reached. If so, the current guess is deemed correct. If not, the current guess is updated using the Jacobian from NIRFAST, and the cycle repeats. The stopping criterion is given by: $\left|\left(\Phi^M - \Phi^C_{(k)}\right)^2\right| > \left|\left(\Phi^M - \Phi^C_{(k-1)}\right)^2\right|$

determined using linear unmixing with their known absorption spectra.

4.2.4 NIRFAST based optical modelling

The update step can be calculated using several methods. In this thesis, the novel approach based on NIRFAST, the so-called virtual detector method, is compared with the direct inversion and the gradient method. This is done using a phantom that contains a large number of small target regions, and one larger known reference region.

As mentioned before, the NIRFAST reconstruction function requires light sources and light detectors, as it is developed for DOT. Light sources are already present in photoacoustic imaging, but instead of optical detectors, "virtual" detectors are placed in the blood where we have determined the Φ^C. This requires a modification to the existing NIRFAST code, as it deletes detectors placed inside the sample (i.e., remove `int_func(~det_boundary_vartex) = 0`). Furthermore, NIRFAST assumes that all source elements can be operated independently of all other sources. In photoacoustic imaging this is often not possible, making it necessary to combine the smaller light sources and everything they affect in the reconstruction steps, such as the fluence and the Jacobian. A limitation of this method is that it is not able to reconstruct the optical scattering coefficient distribution, which is only possible in NIRFAST by determining the direct optical phase shift from pulsed light sources, which cannot be measured with photoacoustic imaging, even with the added reference information. First, the method is showcased on a simple setup with 4 unknown discs and 1 known reference disc. Next, the virtual detector method

is compared in its ability to reconstruct the optical absorption coefficients in photoacoustic imaging simulations with two baseline approaches. Namely, the direct inversion and the gradient method.

Direct inversion, as introduced by Cox *et al.* [36] determines the absorption coefficient distribution by 1) calculating the fluence based on an initial guess of this distribution, 2) dividing the measured absorption coefficient by the calculated fluence, 3) which gives rise to the updated distribution. This iterative approach can be unstable where the calculated fluence is small, which can be dealt with by introducing a constant regularisation factor:

$$\mu_a^{(k+1)} = \hat{H}/(\Phi^{(k)} + \sigma). \tag{4.1}$$

A constant value of 0.01 for σ is used. Equation (4.1) is used as an update step in Figure 4.2 for direct inversion, instead of applying the update step based on the NIRFAST Jacobian. This method requires the optical scattering coefficient distribution to be known.

Next, the gradient method, introduced by Cox *et al.* is discussed [38]. This method sets the sum of squared error between the calculated and measured absorbed energy as the loss function, and utilises a functional gradient for the absorption coefficient distribution based on this squared error:

$$\frac{\partial \mathcal{E}}{\partial \mu_a} = -\Phi \left(\hat{H} - H\right) + \Phi\Phi^*, \tag{4.2}$$

where

$$\left(\mu_a + \nabla \cdot \kappa \nabla\right) \Phi^* = \mu_a \left(\hat{H} - H\right) \tag{4.3}$$

This method is used in combination with `fminunc()` from MATLAB R2022b. The paper of Cox *et al.* also introduced a functional gradient for the reduced scattering coefficient distribution (see Equation (3.17)). However, to make the comparison between the proposed, direct inversion and gradient method equal, only the absorption coefficient gradient is used, and the optical reduced scattering coefficient distribution is assumed to be known for all methods.

After comparing the baseline methods with the novel method, the virtual detector method is examined further by doing a parameter sweep in a simple setup, with one unknown and one known region. The influence of the depth of the test region will be examined, just as the absorption coefficient of the reference region and of the background, and the signal-to-noise ratio. Finally, the novel method is showcased on a phantom containing a blood vessel with a plaque, with tissue-mimicking optical properties.

4.3 Results

4.3.1 Step-by-step explanation of the proposed method

First, the novel method is applied to a simple setup with four unknown discs surrounding one known reference disc on an unknown background, to demonstrate its workings. The μ_a distribution of the phantom can be seen in Figure 4.3A, the unknown discs have an absorption coefficient of 0.05 mm^{-1}, the known reference disc 0.1 mm^{-1}, and the background 0.001 mm^{-1}. The reduced scattering coefficient is 1 mm^{-1} at every location. In Figure 4.3B the setup (sources and detectors) can be seen overlaid on the calculated fluence for the ground truth. This calculated fluence forms the Φ^M are plotted in Figure 4.3C. Φ^C based on the initial guess is also plotted, which is in this case $\mu_a(\mathbf{x}) = 1e-20$, explaining why Φ^C

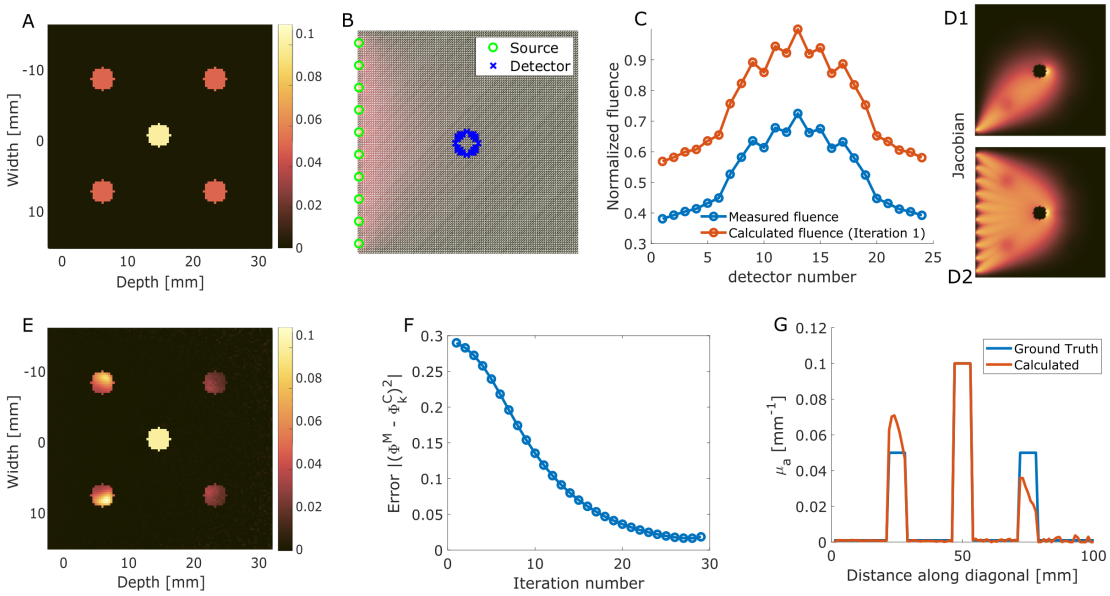

FIGURE 4.3: Demonstration of the proposed method. A: μ_a distribution ground truth, B: calculated fluence overlaid on the computational grid with source and detector location, C: comparison of measured fluence Φ^M and calculated fluence Φ^C for the initial guess, D: one row of the Jacobian matrix corresponding to source 1 and detector 24, and all sources and detector 24, E: reconstructed optical absorption, F: error at each iteration, G: line profile along the diagonal of the phantom comparing the actual and the reconstructed absorption coefficient.

is higher than Φ^M. As mentioned, NIRFAST is based on source-detector pairs. Since it is not possible to measure the fluence from each source point separately, the Jacobian rows for each source-detector pair are merged based on the detector they belong to. One Jacobian row can be seen in Figure 4.3D1 and a combination of several Jacobian rows can be seen in Figure 4.3D2. This method iterates until the error stabilises and starts increasing, where the error per iteration can be seen in 4.3F and the resulting μ_a distribution in Figure 4.3E. To compare the ground truth with the reconstruction output, Figure 4.3G shows an intensity profile along the first diagonal. While the shape of the distribution is comparable (the background is kept low), there is a clear difference in magnitude of the the unknown discs. The maximum reconstructed value is 0.110 mm^{-1}, and the minimum is 0.015 mm^{-1}. However, the average reconstruction value in all the discs is 0.043 ± 0.023 mm^{-1} where the ground truth is 0.050 mm^{-1}.

4.3.2 Proposed vs baseline methods

Having showcased the novel virtual detector method on a simple phantom, we will now compare its performance against two baseline methods: direct inversion and the gradient method. These methods do not require a known reference μ_a as they are not based on virtual fluence detectors. Therefore, we compare the proposed method with four other methods (direct inversion and gradient, with and without reference μ_a as a prior). Figure 4.4 shows their results on a somewhat convoluted phantom, and its quantification is shown in Table 4.1. This phantom contains a known reference disc in the centre (with an absorption coefficient of 0.1 mm^{-1}), surrounded by 18 unknown discs with two different magnitudes (with an absorption coefficient of 0.025 mm^{-1} and 0.05 mm^{-1}) on a back-

ground (with an absorption coefficient of 0.005 mm^{-1}). The reduced scattering coefficient is 1 mm^{-1} at every point.

The first thing that stands out is the fact that for the methods where the reference μ_a is not given as a prior, the direct inversion is not able to accurately reconstruct the reference μ_a, with an average of 0.055 ± 0.005 mm^{-1}, but the gradient method is (0.100 ± 0.001 mm^{-1}). We can compare the other predicted values in their accuracy, where we define accuracy using the percentual error between expected ν_E and average predicted value ν_P: $\frac{\nu_P - \nu_E}{\nu_E} \cdot 100\%$, being larger than 0% for overestimation and below for underestimation. There are two main things (apart from the ground truth and the reconstruction method) that influence this accuracy. Those are the depth at which the unknown disc lies, and the distance between the unknown and reference disc. Looking at the unknown discs that lie at the most superficial depth and relatively close to the reference disc (one of the optimal positions, 1: D1C2), we can see that all methods are able to reconstruct it within 7% absolute error except for the proposed method, which overestimates it with an error of +109%. For the disc that lies somewhat deeper, but at the smallest distance from the reference disc (4: D2C1), we see that other methods dip in accuracy, e.g. direct inversion with an error of -17.4%, while the proposed method improves in accuracy, going from +109% to +64.1% error. Going even further to deeper lying discs on the same width coordinate (12: D4C1 and 15: D5C2), we see that the direct inversion quickly falters, going to errors from -50% till -75%. For these discs, the gradient method with reference known has an error of +11.8% and -46.6% respectively, a smaller drop in accuracy. For the unknown discs lying at the middlemost depth, the proposed method reached on average an error of -32.5% for the discs closest to the reference, and -65.8% for the discs further away from the reference disc. Comparing the ground truth values, for the low-intensity discs, the direct inversion without the reference as a prior has an error of -37.3% on average, while the direct inversion with reference as a prior has an error of -39.4%. For the gradient method, this is -13.7% and +0.6% on average respectively. The proposed method has an average error of -23.7% for these low-intensity discs. For the high-intensity discs, the direct inversion goes from -38.8% to -40.8%, and the gradient method from -7.6% to +2.2%. The proposed method has an average error of -17.9% for these high-intensity discs.

Comparing the standard deviations, which can be quite large for some methods, a clear trend can be observed. The ratio between the standard deviation and the average value that is reconstructed for the discs (the coefficient of variation) gives us an insight into its precision. Ideally, this coefficient is 0. For the four baseline methods, this ratio is below 0.25 for the two most superficial depths and nearly reaches 0.5 for the deepest-lying discs. While there is a clear trend between depth and the coefficient of variation for the baseline methods, there is a dependence on both the depth and the distance to the reference disc for the proposed method. E.g. the highest coefficient of variation is 0.63: one of the deepest lying discs, which is the closest possible to the reference disc (12: D4C1) and the lowest is 0.05: a disc lying as superficial as possible but as far away as possible from the reference disc (2: D1C2$\sqrt{2}$).

Finally, we can compare the direct inversion with and without, and the gradient method with and without the reference as a prior. The direct inversion does not seem to be affected by the inclusion or exclusion of the prior, as the largest difference is between 0.017 mm^{-1} reconstructed with prior and 0.016 mm^{-1}. The exception here is, as mentioned before, the reference disc itself (0.055 ± 0.005 mm^{-1} without prior), but clearly, this does not affect the reconstruction of the other discs. The gradient method, however, seems largely impacted by the prior. For the three most superficial discs, both with and without the prior, the absolute errors stay below 5%, but for the two least superficial depths, they deviate. For

the gradient method with the prior, all three discs at the second deepest level have an error of approximately +10%, whereas the gradient method without the prior has absolute errors below 5%. For the deepest level, this difference becomes even larger, with an error reaching -78% for the gradient method with the prior, while the absolute errors for the gradient method without the prior stay below 20%. In return, the gradient method with the prior has smaller standard deviations than the gradient method without the prior.

Summarising; the accuracy of the direct inversion method does not change with the inclusion of a reference chromophore as a prior, while the accuracy of the gradient method decreases with the inclusion. The direct inversion method drops in accuracy faster than the gradient method when increasing the distance between unknown and illumination, while the opposite is true for the coefficient of variation, both of which are undesirable. The accuracy of the proposed method increases when the distance between the unknown and the reference decreases. Overall, the gradient method without inclusion of the prior consistently achieves the best results.

4.3.3 Characterisation of the proposed method

There are several input parameters which can be changed to characterise the performance of the novel virtual detector method. The selected parameters are the depth of the target disc, the μ_a of the reference disc, the μ_a of the background, and the signal-to-noise ratio. For each parameter, several values are used and for each of these values, it is examined how well the method is able to recover certain μ_a values for the target. The results are shown in Figure 4.5, and quantified in Table 4.2.

As mentioned before, for a changing depth, two effects are influencing the recovery. First, a more superficial (closer to the illumination side) target disc means that more photons are reaching the target, thus the target will have a larger influence on the detected fluence in the reference disc (casting a more intense shadow). Second, a target disc closer to the reference disc means that a larger portion of the photons that could have reached the reference disc is instead absorbed by the target disc, thus having a larger influence on the detected fluence in the reference disc (casting a more precise shadow). These effects are visible in Figure 4.5A, where both the target disc that lies superficial and the target disc that lies close to the reference disc are reconstructed well. For the most superficial case, the error between reconstructed and ground truth is +26.9% for a ground truth of 0.01 mm^{-1}, and almost linearly decreased to +1.7% for 0.1 mm^{-1}. This trend is to a lesser extent also visible for depth=15 mm: going from +33.8% error to -6.8%. However, this trend is not visible for depth=22.5 mm. Instead, it sometimes overestimates and sometimes underestimates. The largest absolute error is for an unknown disc with μ_a of 0.1 mm^{-1}, namely -26.7%. However, it should be noted that the most superficial case reaches a coefficient of variation up to 0.8, while for both other cases, the coefficient stays below 0.5. It is also possible to use the least squares method to calculate a linear fit (of the form $y = a \cdot x + b$) through the average reconstructed values. a will be called the slope, ideally 1, and b the bias, ideally 0. Using this method, both the slope and the bias are the most accurate for the most superficial depth ($a = 1.01$, $b = 0.00$), less accurate for the middle value ($a = 0.93$, $b = 0.01$), and the least accurate for the deepest value ($a = 0.74$, $b = 0.01$).

For a changing reference μ_a, it can be seen that the recovery of the target disc is done accurately (with absolute errors below +20% and coefficients of variation below 0.6), except for the case where the reference μ_a is equal to that of the background. For the latter, an error of -81.3% is reached and a coefficient of variation of 1.1, both for a ground truth of 0.1 mm^{-1}. If we look at the least squared method, the best fit is acquired for a reference

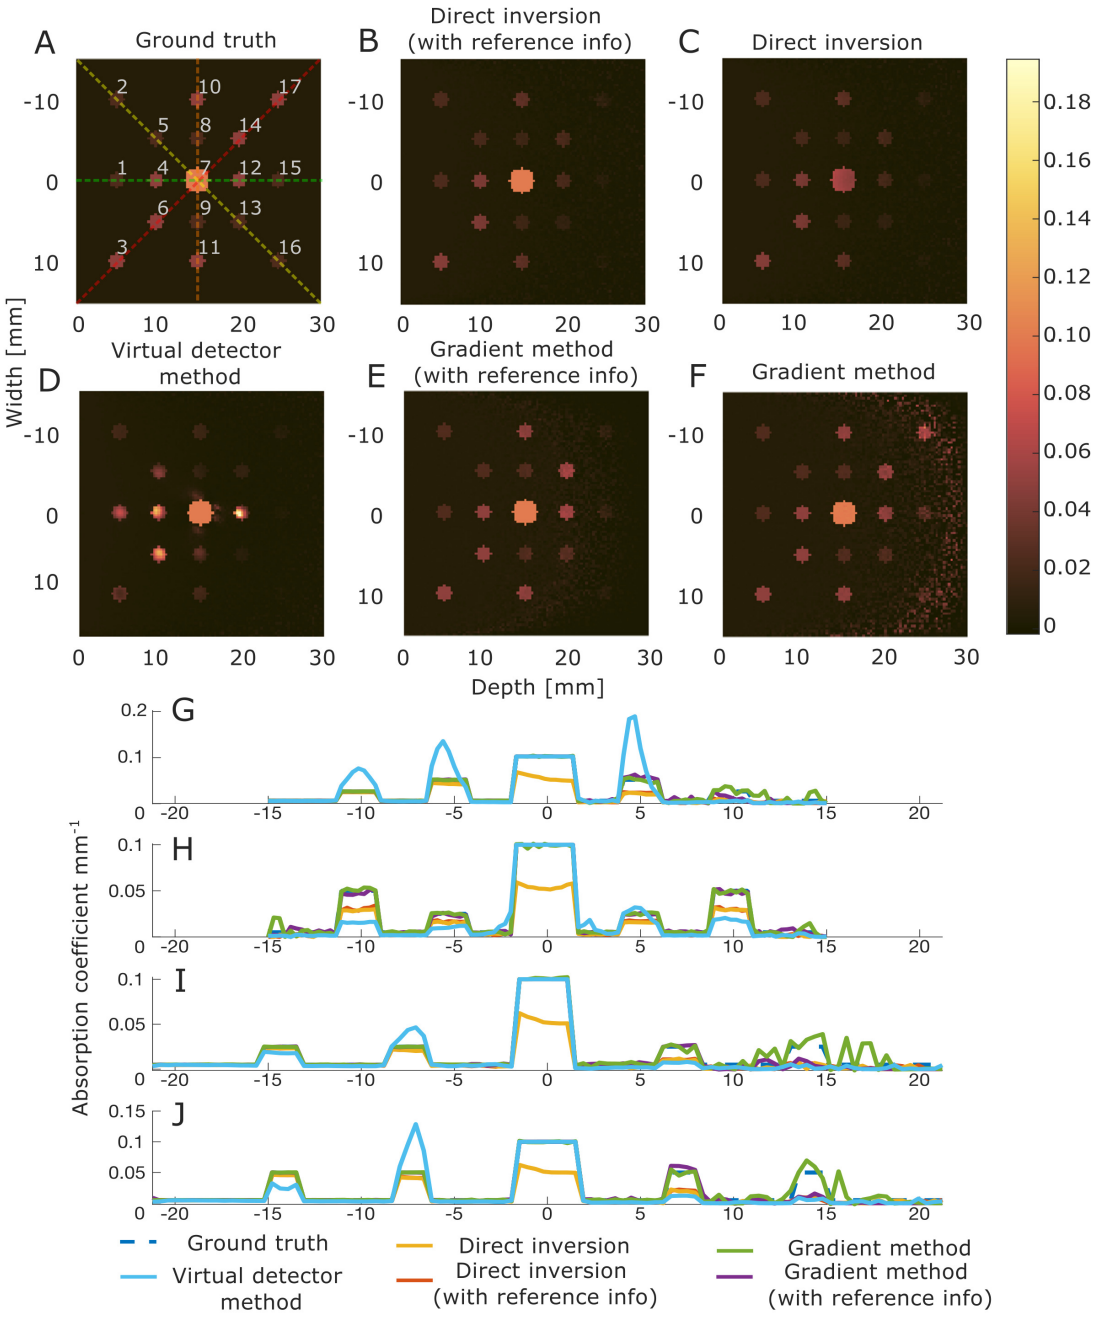

FIGURE 4.4: A-F: Comparison of the output of the different methods: direct inversion (with and without reference μ_a known), gradient (with and without reference μ_a known) and novel method. Illumination from the left side. Also showing intensity line profiles along G: horizontal centre-line (Green), H: vertical centre-line (Orange), I: top left-bottom right diagonal (Yellow), J: top right-bottom left diagonal (Red). Quantification is done in Table 4.1.

Code	Ground truth	Direct inversion with reference known	Direct inversion	Gradient with reference known	Gradient	Proposed method
1: D1C2	0.025	0.024 ± 0.000	0.023 ± 0.000	0.025 ± 0.000	0.025 ± 0.000	0.052 ± 0.011
2: D1C2$\sqrt{2}$	0.025	0.023 ± 0.000	0.023 ± 0.000	0.025 ± 0.000	0.025 ± 0.000	0.019 ± 0.001
3: D1C2$\sqrt{2}$	0.05	0.046 ± 0.000	0.046 ± 0.001	0.050 ± 0.000	0.050 ± 0.000	0.030 ± 0.004
4: D2C1	0.05	0.043 ± 0.001	0.041 ± 0.001	0.050 ± 0.000	0.050 ± 0.000	0.082 ± 0.032
5: D2C$\sqrt{2}$	0.025	0.021 ± 0.000	0.021 ± 0.001	0.025 ± 0.000	0.025 ± 0.000	0.035 ± 0.008
6: D2C$\sqrt{2}$	0.05	0.042 ± 0.001	0.041 ± 0.001	0.050 ± 0.000	0.050 ± 0.000	0.084 ± 0.024
7: D3C0	0.1	0.100 ± 0.000	0.055 ± 0.005	0.100 ± 0.000	0.100 ± 0.001	0.100 ± 0.000
8: D3C1	0.025	0.017 ± 0.001	0.016 ± 0.001	0.025 ± 0.001	0.025 ± 0.001	0.010 ± 0.002
9: D3C1	0.025	0.017 ± 0.001	0.016 ± 0.001	0.025 ± 0.002	0.025 ± 0.001	0.023 ± 0.006
10: D3C2	0.05	0.030 ± 0.002	0.029 ± 0.002	0.048 ± 0.002	0.050 ± 0.002	0.016 ± 0.001
11: D3C2	0.05	0.030 ± 0.002	0.029 ± 0.002	0.049 ± 0.002	0.050 ± 0.001	0.018 ± 0.002
12: D4C1	0.05	0.022 ± 0.001	0.020 ± 0.002	0.056 ± 0.003	0.049 ± 0.003	0.081 ± 0.052
13: D4C$\sqrt{2}$	0.025	0.011 ± 0.001	0.011 ± 0.001	0.028 ± 0.003	0.024 ± 0.003	0.007 ± 0.001
14: D4C$\sqrt{2}$	0.05	0.022 ± 0.002	0.021 ± 0.002	0.055 ± 0.004	0.050 ± 0.004	0.012 ± 0.001
15: D5C2	0.025	0.007 ± 0.002	0.006 ± 0.002	0.013 ± 0.004	0.027 ± 0.006	0.004 ± 0.001
16: D5C2$\sqrt{2}$	0.025	0.005 ± 0.002	0.005 ± 0.002	0.007 ± 0.003	0.024 ± 0.011	0.003 ± 0.002
17: D5C2$\sqrt{2}$	0.05	0.009 ± 0.003	0.009 ± 0.002	0.011 ± 0.003	0.060 ± 0.013	0.006 ± 0.001

TABLE 4.1: Average reconstructed values and their standard deviations, as shown in Figure 4.4. All values are in mm^{-1}, with the exception of the code. The code indicates the location of the disc; D for the depth: 1 meaning most superficial, 5, least superficial; C for distance from reference disc, 1 for closest to reference disc, $2\sqrt{2}$ for furthest. The disc C0 is the reference disc.

absorption coefficient of 0.05 mm^{-1}: $a = 1.00$, $b = 0.00$. For a reference μ_a of 0.001 mm^{-1} these are 0.67 and 0.01 respectively.

For a changing background μ_a, a limitation is quickly reached: for values of 0.005 and 0.01 mm^{-1}, the reconstruction strongly deviates from the ground truth. When the background μ_a is 0.01 mm^{-1} an error of +169.6% is reached (for a ground truth of 0.05 mm^{-1}). For both a background of 0.005 and 0.01 mm^{-1}, the coefficient of variation rises with the rising ground truth, starting with 0.25 for 0.005 mm^{-1} and 0.26 for 0.01 mm^{-1}, and ending with 0.50 for 0.005 mm^{-1} and 1.12 for 0.01 mm^{-1}. For the two other background μ_a cases, the absolute error stays below 15%, and the coefficient of variation below 0.35, with one exception (0.5 for a ground truth of 0.8 mm^{-1} and a background of 0.001 mm^{-1}). The results of the linear regression give a comparable analysis, the slope of both background μ_a of 0.005 and 0.01 mm^{-1} are high: 1.67 and 1.89 respectively. The best result is acquired for 0.001 mm^{-1} ($a = 1.08$, $b = 0.00$).

For a changing SNR, the accuracy of the reconstruction is directly correlated: a higher SNR means a higher accuracy. For the two lowest SNR values that were tested, a constant underestimation is the case. The error for the lowest SNR value stays constantly below -40% with a lowest value of -84.7%, and for the second lowest SNR value below -15% (with one exception) with a lowest value of -55.5%. The other SNR values result in better reconstructions. However, for the case where the SNR is 60 dB, the method starts overestimating (e.g. 0.124 vs 0.1 mm^{-1}). For both the under- and overestimation, the added noise is a likely reason. Since the noise is not allowed to make the absorbed energy map negative, as that would not make physical sense: $\hat{H} = max\,(H + \mathcal{N}, 0)$, the average measured fluence Φ^M could be higher than without noise. Namely, the noise is implemented as white Gaussian noise. A higher measured fluence would mean that less light was absorbed, meaning a lower absorption coefficient. A reason for the overestimation could be that the noise has strongly increased one of the more important detector pixels, therefore hampering correct reconstruction. The coefficient of variation stays below 0.55 for all of the four highest SNR values. A significant detail is that the coefficients of variation in SNR=60 are almost always higher than those for SNR=30 in otherwise comparable settings. This again points to critical noise as a reason. Furthermore, the least squares method to calculate a linear fit also indicates the over- and underestimation. For an SNR of 5 and 10, the slope is below 1: 0.49 and 0.61 respectively, for an SNR of 60 the slope is above 1: 1.19. For all SNR values the absolute value of the bias is low, always below 0.02.

4.3.4 Application of the proposed method on plaque phantom

Having characterised the proposed method, its abilities are now showcased on a phantom which approximates a realistic setup: a muscle layer (platysma muscle), a muscle half ellipse (sternocleidomastoid muscle), and a blood disc (Ox.stat. of 95%) with a plaque (90% lipid, 10% collagen). The phantom is shown in Figure 4.7A. Both muscle areas, the blood and the plaque have correct optical proprieties for an optical wavelength of 800 nm [50, 51, 52]. Again, the proposed virtual detector method is applied by first determining the absorbed energy map through simulation with NIRFAST. This absorbed energy map is then used to determine the measured fluence Φ^M at the virtual detectors, which the proposed method tries to approach by iteratively updating the initial guess of the absorption coefficient distribution. The reduced scattering coefficient distribution is assumed to be known. The reconstructed absorption coefficient distribution can be seen in Figure 4.7B. While both muscle areas show large deviations from the correct values, the plaque reconstruction has better accuracy. The blood optical properties are of course given as a prior and thus match flawlessly. The superficial muscle layer has an average reconstructed value of $0.13 \pm$

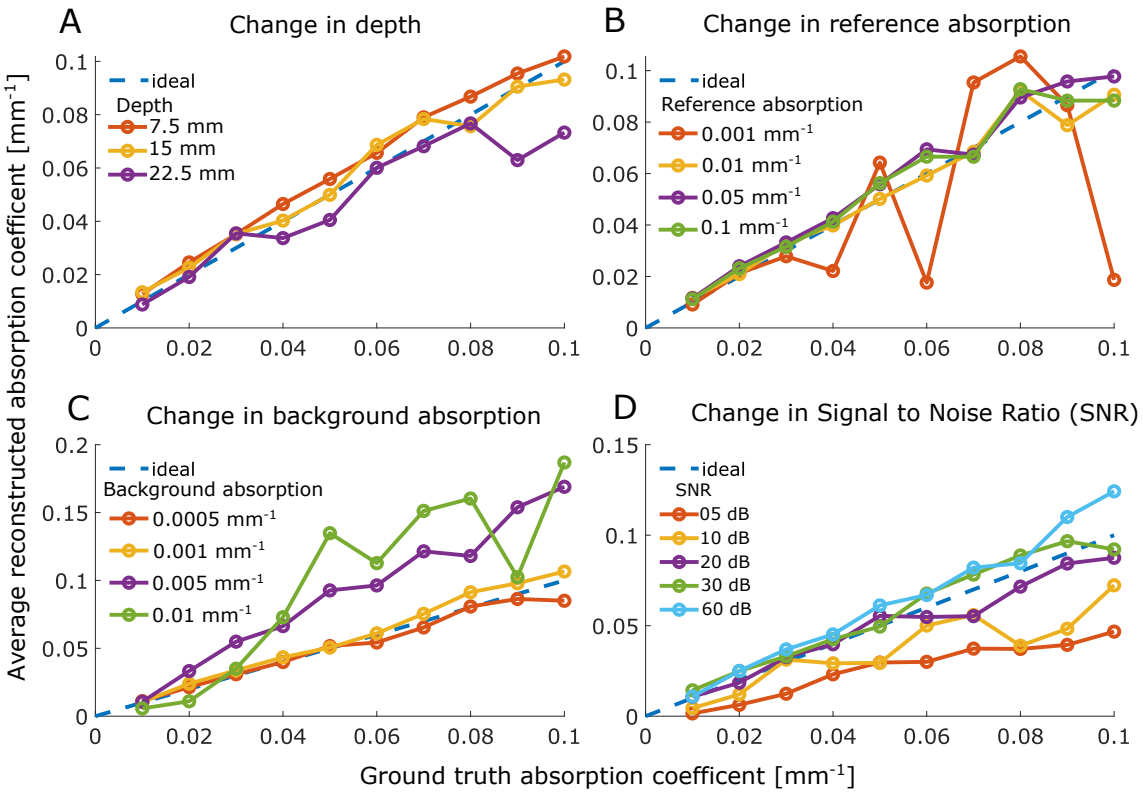

FIGURE 4.5: Recovery of the target disc μ_a, while changing several parameters: A: depth of the test disc as indicated in Figure 4.6, B: reference disc μ_a, C: background μ_a, D: signal-to-noise ratio (SNR). Normal values are depth=15 mm (see Figure 4.6), reference disc $\mu_a = 0.01$ mm^{-1}, background $\mu_a = 0.001$ mm^{-1}, SNR=30. Quantification is done in Table 4.2.

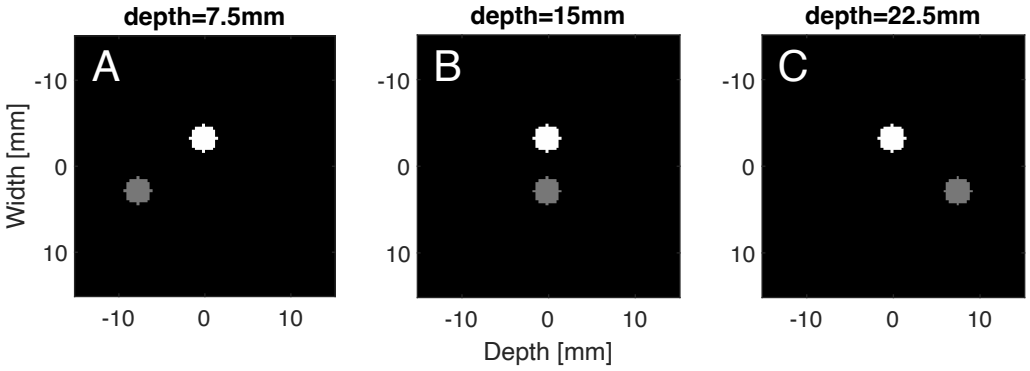

FIGURE 4.6: Changing depth of the reference disc, as used in Figure 4.5. The target disc has a depth of: A: 7.5mm, B: 15mm, C: 22.5mm. The reference disc always has a depth of 15mm, illumination from the left side. Illumination from the left side.

	Depth			SNR				
Ground truth	7.5 mm	15 mm	22.5 mm	5	10	20	30	60
0.01	0.013 ± 0.001	0.013 ± 0.002	0.009 ± 0.001	0.002 ± 0.001	0.004 ± 0.001	0.011 ± 0.001	0.014 ± 0.002	0.011 ± 0.005
0.02	0.025 ± 0.004	0.023 ± 0.007	0.019 ± 0.004	0.006 ± 0.002	0.012 ± 0.003	0.019 ± 0.004	0.025 ± 0.007	0.025 ± 0.007
0.03	0.035 ± 0.006	0.035 ± 0.008	0.035 ± 0.011	0.012 ± 0.004	0.031 ± 0.005	0.033 ± 0.010	0.033 ± 0.010	0.037 ± 0.013
0.04	0.046 ± 0.006	0.040 ± 0.011	0.034 ± 0.009	0.023 ± 0.010	0.029 ± 0.016	0.040 ± 0.007	0.043 ± 0.007	0.045 ± 0.014
0.05	0.056 ± 0.023	0.050 ± 0.015	0.041 ± 0.013	0.030 ± 0.011	0.029 ± 0.002	0.055 ± 0.010	0.049 ± 0.010	0.061 ± 0.027
0.06	0.066 ± 0.016	0.069 ± 0.022	0.060 ± 0.024	0.030 ± 0.021	0.050 ± 0.007	0.055 ± 0.018	0.068 ± 0.020	0.067 ± 0.013
0.07	0.079 ± 0.033	0.078 ± 0.023	0.068 ± 0.026	0.037 ± 0.005	0.056 ± 0.017	0.055 ± 0.029	0.078 ± 0.015	0.082 ± 0.029
0.08	0.087 ± 0.030	0.076 ± 0.028	0.077 ± 0.033	0.037 ± 0.018	0.039 ± 0.015	0.072 ± 0.029	0.089 ± 0.043	0.085 ± 0.017
0.09	0.095 ± 0.025	0.091 ± 0.016	0.063 ± 0.018	0.039 ± 0.010	0.048 ± 0.019	0.084 ± 0.017	0.097 ± 0.030	0.110 ± 0.049
0.1	0.102 ± 0.081	0.093 ± 0.027	0.073 ± 0.029	0.047 ± 0.016	0.072 ± 0.017	0.088 ± 0.025	0.092 ± 0.027	0.124 ± 0.063
	Reference μ_a				Background μ_a			
Ground truth	0.001 mm^{-1}	0.01 mm^{-1}	0.05 mm^{-1}	0.1 mm^{-1}	0.0005 mm^{-1}	0.001 mm^{-1}	0.005 mm^{-1}	0.01 mm^{-1}
0.01	0.009 ± 0.001	0.012 ± 0.001	0.012 ± 0.001	0.011 ± 0.001	0.011 ± 0.001	0.011 ± 0.002	0.011 ± 0.003	0.006 ± 0.001
0.02	0.021 ± 0.009	0.021 ± 0.005	0.024 ± 0.004	0.023 ± 0.007	0.021 ± 0.007	0.024 ± 0.005	0.033 ± 0.010	0.011 ± 0.002
0.03	0.028 ± 0.012	0.033 ± 0.008	0.033 ± 0.008	0.032 ± 0.008	0.031 ± 0.004	0.034 ± 0.010	0.055 ± 0.022	0.035 ± 0.015
0.04	0.022 ± 0.007	0.040 ± 0.006	0.043 ± 0.009	0.041 ± 0.010	0.040 ± 0.014	0.043 ± 0.010	0.066 ± 0.030	0.073 ± 0.041
0.05	0.064 ± 0.053	0.050 ± 0.026	0.056 ± 0.022	0.056 ± 0.012	0.052 ± 0.008	0.051 ± 0.014	0.093 ± 0.040	0.135 ± 0.087
0.06	0.018 ± 0.012	0.059 ± 0.021	0.069 ± 0.021	0.067 ± 0.014	0.054 ± 0.014	0.061 ± 0.015	0.096 ± 0.034	0.113 ± 0.084
0.07	0.095 ± 0.068	0.068 ± 0.042	0.067 ± 0.012	0.067 ± 0.020	0.065 ± 0.018	0.075 ± 0.021	0.121 ± 0.050	0.151 ± 0.127
0.08	0.106 ± 0.031	0.092 ± 0.034	0.090 ± 0.024	0.093 ± 0.042	0.081 ± 0.017	0.091 ± 0.049	0.118 ± 0.051	0.160 ± 0.104
0.09	0.087 ± 0.019	0.079 ± 0.029	0.096 ± 0.030	0.088 ± 0.018	0.086 ± 0.019	0.098 ± 0.017	0.154 ± 0.078	0.102 ± 0.066
0.1	0.019 ± 0.021	0.091 ± 0.045	0.098 ± 0.014	0.088 ± 0.035	0.085 ± 0.022	0.107 ± 0.035	0.169 ± 0.085	0.187 ± 0.209

TABLE 4.2: Average reconstructed values and their standard deviations, as shown in Figure 4.5. All values are in mm^{-1}, with the exception of the SNR.

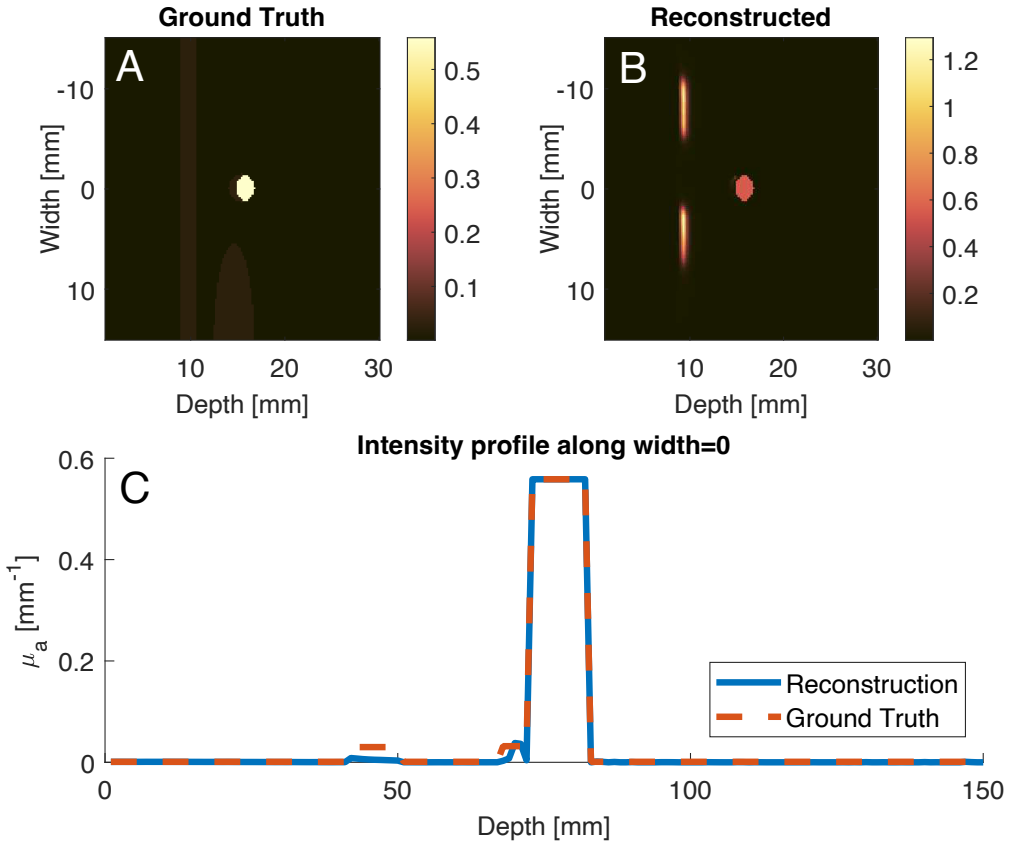

FIGURE 4.7: Application of the proposed method on a phantom with accurate anatomical features and optical properties. A: ground truth of μ_a distribution, B: reconstructed μ_a distribution, C: intensity profile of ground truth and reconstruction along width centre-line. Values used: muscle areas: μ_a: 0.03 mm^{-1}, μ_s': 0.6 mm^{-1}; blood: μ_a: 0.559 mm^{-1}, μ_s': 1.31 mm^{-1}; plaque: μ_a: 0.032 mm^{-1}, μ_s': 13.2 mm^{-1}; background: μ_a: 0.001 mm^{-1}, μ_s': 1 mm^{-1}.

0.26 mm^{-1} (thus a coefficient of variation of 2.0) and the deeper muscle half ellipse an average reconstructed value of 0.0005 ± 0.0003 mm^{-1} (coefficient of variation of 0.7), where the correct value is 0.03 mm^{-1}. This means an error of +338.25% and -98.32% respectively. More interestingly, the reconstructed value of the plaque is 0.034 ± 0.044 mm^{-1}, and the correct value is 0.03 mm^{-1}. This results in a large coefficient of variance of 1.28, but a relatively small average error of +6.96%.

Chapter 5

Conclusions and Future Work

In this thesis, an algorithm was created that is able to utilise a fluence marker during the reconstruction and quantification of tissue composition based on photoacoustic signals. For this goal, two separate research questions were formulated, which will now be treated in turn.

BME RQ: Can the integration of an endogenous reference chromophore, such as blood in the carotid artery, enhance the accuracy of quantitative photoacoustic imaging in the retrieval of non-reference chromophore concentrations, such as in plaques in the carotid artery?

Starting with the BME RQ, no clear improvement in accuracy can be seen when including a reference chromophore as an optical prior. As highlighted in Section 4.3.2, the direct inversion method seems largely unaffected by the inclusion of the prior. However, the gradient method even decreases in accuracy for deeper lying unknown discs compared to the same method without the inclusion of the prior (approx. 10%pt. up to 70%pt. absolute error increase). During the direct inversion, each pixel is relatively unaffected by the optical values of the other pixels: each pixel is inverted separately. The only connection between them is that the current values of the surrounding pixels influence the calculated fluence for the pixel, which is used to calculate the subsequent value. In the gradient method, this connection is stronger, as the current value of the surrounding pixels determines the calculated fluence, the calculated absorbed energy distribution, and the adjoint operator (from Equation (3.18)). Therefore, it is not remarkable that the direct inversion is largely unaffected by the exclusion of the prior, but the gradient method is affected. However, the gradient method decreases when one includes the prior, instead of increasing in accuracy. This is unexpected, as one could expect that including more prior knowledge increases any reconstruction algorithm. An explanation for this could be the implementation of the prior in the gradient method. Currently, the gradient method is implemented using the `fminunc()` function from MATLAB. This algorithm tries to find the minimum of an unconstrained function, which here is Equation (3.15), with Equation (3.16) supplied as a gradient. As can be expected based on the name of the function, it is unconstrained and can therefore not directly handle constraints. The way the constraints are implemented is by providing and overwriting the reference values that are used in the calculation of the gradient and error. However, the minimisation function has no knowledge of these reference values: its internal values are not overwritten and can and will therefore change these values without effect. This might be the cause for the minimisation function not reaching accurate values. There is a comparable function in MATLAB that is able to handle constraints: `fmincon()`, but defining strict boundaries for the values the reference disc is able to take will cause problems for its default algorithm. Furthermore,

this constrained function is multiple orders of magnitude more computationally demanding. Another option would be to first run the reconstruction without the prior, and then use the result as an initial guess for reconstruction with prior. A final note on this research question is that while an increase in accuracy cannot be found in the phantom used in Figure 4.4A, the increase might only be visible if the separation between reference and unknown is even smaller. The proposed virtual detector method (see next paragraph) does not outperform the other methods in accuracy, but when applying it on a plaque phantom where the reference chromophore lies directly next to the target, a small error of approximately +7% is reached. However, the standard deviation is quite large, indicating that the reconstruction is not uniform.

CS RQ: Can the incorporation of virtual detectors (requiring certain optical properties and spatial location as prior) enhance the accuracy of model-based optical inversion in quantitative photoacoustic imaging?

Moving on with the CS RQ, virtual detectors were incorporated in model-based optical inversion. The reconstruction engine of NIRFAST was used as a starting point, which requires multiple light sources and multiple light detectors at the tissue boundary to reconstruct the optical properties. By modifying NIRFAST to allow detectors inside the tissue, instead of only at the boundary, it is possible to use a reference chromophore as a fluence marker. Furthermore, NIRFAST determines the fluence and a Jacobian row for each source-detector combination, which was modified to such that it determines those for each detector (thus aggregating the sources for each detector). Together, these modifications make it possible to use the reconstruction engine from NIRFAST for quantitative photoacoustic imaging with virtual detectors based on a reference chromophore, i.e. having its optical properties and spatial location as a prior. With this modified engine, results within 20% absolute error can be achieved, and for optimal cases within 10% absolute error. The method is very dependent on the distance between the reference and target: a large distance (e.g., > 5 mm) can dampen its accuracy quickly, as expected. It was also found that the reference absorption coefficient should be high (≥ 0.01 mm^{-1}), while the background absorption coefficient should be low (≤ 0.001 mm^{-1}). For the latter, it is expected that a high accuracy can still be reached for higher values, as long as the distances between illumination, reference and target are smaller. Namely, a higher background absorption coefficient will result in a lower fluence and lower PA signal in the reference and the target, which can be compensated by reducing the distance between illumination, reference and target. Furthermore, the SNR should be higher than 10 dB, lower and its absolute error in the tested setting was larger than 40%. While the average error can be quite low, the standard deviation is often large, pointing to a non-uniform result. This can be mitigated in several ways, one of which is the usage of soft spatial priors. As described by Davis *et al.* [53], it is possible to include spatial priors in the cost function that is used by NIRFAST. This modification in the form of a Laplacian-type structure has the goal to minimise variability within a region, but allow gradual changes between different regions in directions normal to their common boundary. Instead of soft spatial priors, one can also have hard spatial priors. By applying the same technique used to combine the Jacobian rows for source-detector combinations, one can also combine the Jacobian columns such that each region is updated uniformly. This means that each region is treated as one very large pixel when it comes to calculating and applying the update step. Where the soft spatial priors method allows variability within a region, with hard spatial priors all values in one region are and stay the same. Of course, these spatial priors need to be acquired. This can relatively easily be done using US, as the PA setup already includes an acoustic transducer. By applying segmentation techniques on the B-mode images of the US trans-

ducer, one can split the imaging plane into separate tissue types. Another option is MRI, but this requires separate imaging and is more costly.

As mentioned, future work can focus on better integrating the reference chromophore into the existing baseline methods (e.g. using `fmincon()`) or improve the virtual detector method (by including spatial priors). A different improvement that can be made to the proposed method with NIRFAST stems from the fact that it is currently only reconstructing for a specific optical wavelength. It is possible to extend this reconstruction step to several wavelengths at once. This means introducing constraints based on the absorption spectra of the present chromophores and including the concentration step in the reconstruction. Furthermore, it is also possible to use data-driven techniques instead of model-based approaches. For example, physics-based deep learning approaches can even exploit our knowledge of light propagation during the prediction of optical properties [54]. Next, while the optical inversion method using a modified method might eventually work very well, it should also be able to handle the distortions that are introduced by the acoustic inversion. This acoustic inversion was assumed to be done accurately but should be included in the pipeline from measurement to concentration, and the eventual problems that arise in the optical inversion because of the non-ideal acoustic inversion should be mitigated. Finally, the analysis and conclusions made in this thesis are purely based on (optical) simulations. Although these might translate well into practice, it is best to verify them in an experimental setting. This might require improving the robustness of the method somewhat, but one should be wary of optimising the method purely based on simulations for too long.

In summary, it is possible to include a reference chromophore in quantitative photoacoustic imaging, by integrating it into existing methods, or by using it with a modified method from DOT. The latter, called the virtual detector method, is capable of reconstructing optical properties in photoacoustic imaging through the usage of virtual detectors in the reference chromophore. This technique might, with some improvement, be used to accurately quantify the optical properties of a plaque, with blood as the reference chromophore. By repeating this for several optical wavelengths, one can then determine the composition of the plaque and which is directly connected to its stability. This information will help reduce over- and undertreatment.

Acknowledgement

I would like to thank Francis Kalloor Joseph and Anjali Thomas for supporting and helping me the past year, and my graduation committee for their expertise. I am also grateful to Nienke Bosschaart, Jelmer Wolterink and Wiendelt Steenbergen for their input during my work. Many thanks to Hamid Dehghani and Jiaming Cao for their help with NIRFAST. Finally, I would like to thank the people from the BMPI group, as well as Emiel, my friends, and my family for their (moral) support.

Bibliography

[1] J. F. Meschia, C. Bushnell, B. Boden-Albala, L. T. Braun, D. M. Bravata, S. Chaturvedi, M. A. Creager, R. H. Eckel, M. S. Elkind, M. Fornage, L. B. Goldstein, S. M. Greenberg, S. E. Horvath, C. Iadecola, E. C. Jauch, W. S. Moore, and J. A. Wilson, "Guidelines for the primary prevention of stroke," *Stroke*, vol. 45, no. 12, pp. 3754–3832, 2014.

[2] H. J. M. Barnett, R. W. Gunton, M. Eliasziw, L. Fleming, B. Sharpe, P. Gates, and H. Meldrum, "Causes and Severity of Ischemic Stroke in Patients With Internal Carotid Artery Stenosis," *JAMA*, vol. 283, no. 11, pp. 1429–1436, 2000.

[3] F. Sharif, D. Lohan, and W. Wijns, "Non-invasive detection of vulnerable coronary plaque," *World journal of cardiology*, vol. 3, pp. 219–29, 2011.

[4] C. Oates, A. Naylor, T. Hartshorne, S. Charles, T. Fail, K. Humphries, M. Aslam, and P. Khodabakhsh, "Joint recommendations for reporting carotid ultrasound investigations in the united kingdom," *European Journal of Vascular and Endovascular Surgery*, vol. 37, no. 3, pp. 251–261, 2009.

[5] B. T. Cox, J. G. Laufer, P. C. Beard, and S. R. Arridge, "Quantitative spectroscopic photoacoustic imaging: a review," *Journal of Biomedical Optics*, vol. 17, no. 6, p. 061202, 2012.

[6] J. L. Sandell and T. C. Zhu, "A review of in-vivo optical properties of human tissues and its impact on pdt," *Journal of Biophotonics*, vol. 4, no. 11-12, pp. 773–787, 2011.

[7] Y. Yuan, S. Yan, and Q. Fang, "Light transport modeling in highly complex tissues using the implicit mesh-based monte carlo algorithm," *Biomedical Optics Express*, vol. 12, no. 1, pp. 147–161, 2021.

[8] G. M. Palmer and N. Ramanujam, "Monte carlo-based inverse model for calculating tissue optical properties. part i: Theory and validation on synthetic phantoms," *Applied Optics*, vol. 45, no. 5, pp. 1062–1071, 2006.

[9] J. Buchmann, B. A. Kaplan, S. Powell, S. Prohaska, and J. Laufer, "Three-dimensional quantitative photoacoustic tomography using an adjoint radiance Monte Carlo model and gradient descent," *Journal of Biomedical Optics*, vol. 24, no. 6, p. 066001, 2019.

[10] S. Manohar and D. Razansky, "Photoacoustics: a historical review," *Advances in Optics and Photonics*, vol. 8, no. 4, pp. 586–617, 2016.

[11] P. Beard, "Biomedical photoacoustic imaging," *Interface Focus*, vol. 1, no. 4, pp. 602–631, 2011.

[12] M. Xu and L. V. Wang, "Universal back-projection algorithm for photoacoustic computed tomography," *Physical Review E*, vol. 71, p. 016706, 2005.

[13] A. B. Treeby, B. Cox, and J. Jaros, "k-wave: A matlab toolbox for the time domain simulation of acoustic wave fields, user manual."

[14] B. Fornberg, "The pseudospectral method; comparisons with finite differences for the elastic wave equation," *Geophysics*, vol. 52, no. 4, pp. 483–501, 1987.

[15] B. E. Treeby, J. Jaros, A. P. Rendell, and B. T. Cox, "Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using a k-space pseudospectral method," *The Journal of the Acoustical Society of America*, vol. 131, no. 6, pp. 4324–4336, 2012.

[16] R. A. Kruger, R. B. Lam, D. R. Reinecke, S. P. Del Rio, and R. P. Doyle, "Photoacoustic angiography of the breast," *Medical Physics*, vol. 37, no. 11, pp. 6096–6100, 2010.

[17] Y. Zhang and Y. Wang, "An improved filtered back-projection algorithm for photoacoustic tomography," in *2011 5th International Conference on Bioinformatics and Biomedical Engineering*, pp. 1–4, 2011.

[18] M. Mozaffarzadeh, A. Mahloojifar, M. Orooji, S. Adabi, and M. Nasiriavanaki, "Double-stage delay multiply and sum beamforming algorithm: Application to linear-array photoacoustic imaging," *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 1, pp. 31–42, 2018.

[19] A. M. Deylami and B. M. Asl, "Low complex subspace minimum variance beamformer for medical ultrasound imaging," *Ultrasonics*, vol. 66, pp. 43–53, 2016.

[20] G. A. Hendriks, C. Chen, H. H. Hansen, and C. L. De Korte, "Quasi-static elastography and ultrasound plane-wave imaging: The effect of beam-forming strategies on the accuracy of displacement estimations," *Applied Sciences*, vol. 8, no. 3, 2018.

[21] B. E. Treeby and B. T. Cox, "k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields," *Journal of Biomedical Optics*, vol. 15, no. 2, p. 021314, 2010.

[22] S. R. Arridge, M. M. Betcke, B. T. Cox, F. Lucka, and B. E. Treeby, "On the adjoint operator in photoacoustic tomography," *Inverse Problems*, vol. 32, no. 11, p. 115012, 2016.

[23] K. Wang, S. A. Ermilov, R. Su, H.-P. Brecht, A. A. Oraevsky, and M. A. Anastasio, "An imaging model incorporating ultrasonic transducer properties for three-dimensional optoacoustic tomography," *IEEE Transactions on Medical Imaging*, vol. 30, no. 2, pp. 203–214, 2011.

[24] D. Van de Sompel, L. S. Sasportas, J. V. Jokerst, and S. S. Gambhir, "Comparison of deconvolution filters for photoacoustic tomography," *PLOS ONE*, vol. 11, no. 3, pp. 1–28, 2016.

[25] J. H. Friedman, T. Hastie, and R. Tibshirani, "Regularization paths for generalized linear models via coordinate descent," *Journal of Statistical Software*, vol. 33, no. 1, p. 1–22, 2010.

[26] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[27] J. Prakash, A. S. Raju, C. B. Shaw, M. Pramanik, and P. K. Yalavarthy, "Basis pursuit deconvolution for improving model-based reconstructed images in photoacoustic tomography," *Biomedical Optics Express*, vol. 5, no. 5, pp. 1363–1377, 2014.

[28] J. Schwab, S. Antholzer, and M. Haltmeier, "Learned backprojection for sparse and limited view photoacoustic tomography," in *Photons Plus Ultrasound: Imaging and Sensing 2019* (A. A. Oraevsky and L. V. Wang, eds.), vol. 10878, p. 1087837, International Society for Optics and Photonics, SPIE, 2019.

[29] S. Antholzer, J. Schwab, J. Bauer-Marschallinger, P. Burgholzer, and M. Haltmeier, "NETT regularization for compressed sensing photoacoustic tomography," in *Photons Plus Ultrasound: Imaging and Sensing 2019* (A. A. Oraevsky and L. V. Wang, eds.), vol. 10878, p. 108783B, International Society for Optics and Photonics, SPIE, 2019.

[30] A. Hauptmann, F. Lucka, M. Betcke, N. Huynh, J. Adler, B. Cox, P. Beard, S. Ourselin, and S. Arridge, "Model-based learning for accelerated, limited-view 3-d photoacoustic tomography," *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1382–1393, 2018.

[31] E. M. A. Anas, H. K. Zhang, C. Audigier, and E. M. Boctor, "Robust photoacoustic beamforming using dense convolutional neural networks," in *Simulation, Image Processing, and Ultrasound Systems for Assisted Diagnosis and Navigation* (D. Stoyanov, Z. Taylor, S. Aylward, J. M. R. Tavares, Y. Xiao, A. Simpson, A. Martel, L. Maier-Hein, S. Li, H. Rivaz, I. Reinertsen, M. Chabanas, and K. Farahani, eds.), (Cham), pp. 3–11, Springer International Publishing, 2018.

[32] D. Waibel, J. Gröhl, F. Isensee, T. Kirchner, K. Maier-Hein, and L. Maier-Hein, "Reconstruction of initial pressure from limited view photoacoustic images using deep learning," in *Photons Plus Ultrasound: Imaging and Sensing 2018* (A. A. Oraevsky and L. V. Wang, eds.), vol. 10494, p. 104942S, International Society for Optics and Photonics, SPIE, 2018.

[33] M. Kim, G.-S. Jeng, I. Pelivanov, and M. O'Donnell, "Deep-learning image reconstruction for real-time photoacoustic system," *IEEE Transactions on Medical Imaging*, vol. 39, no. 11, pp. 3379–3390, 2020.

[34] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems* (Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, eds.), vol. 27, Curran Associates, Inc., 2014.

[35] J. Gröhl, M. Schellenberg, K. Dreher, and L. Maier-Hein, "Deep learning for biomedical photoacoustic imaging: A review," *Photoacoustics*, vol. 22, p. 100241, 2021.

[36] B. T. Cox, S. R. Arridge, K. P. Köstli, and P. C. Beard, "Two-dimensional quantitative photoacoustic image reconstruction of absorption distributions in scattering media by use of a simple iterative method," *Applied Optics*, vol. 45, no. 8, pp. 1866–1875, 2006.

[37] B. Banerjee, S. Bagchi, R. M. Vasu, and D. Roy, "Quantitative photoacoustic tomography from boundary pressure measurements: noniterative recovery of optical absorption coefficient from the reconstructed absorbed energy map," *Journal of the Optical Society of America A*, vol. 25, no. 9, pp. 2347–2356, 2008.

[38] B. T. Cox, S. R. Arridge, and P. C. Beard, "Gradient-based quantitative photoacoustic image reconstruction for molecular imaging," in *Photons Plus Ultrasound: Imaging and Sensing 2007: The Eighth Conference on Biomedical Thermoacoustics, Optoacoustics, and Acousto-optics* (A. A. Oraevsky and L. V. Wang, eds.), vol. 6437, p. 64371T, International Society for Optics and Photonics, SPIE, 2007.

[39] C. Cai, K. Deng, C. Ma, and J. Luo, "End-to-end deep neural network for optical inversion in quantitative photoacoustic imaging," *Optics Letters*, vol. 43, no. 12, pp. 2752–2755, 2018.

[40] J. Gröhl, T. Kirchner, T. Adler, and L. Maier-Hein, "Confidence estimation for machine learning-based quantitative photoacoustics," *Journal of Imaging*, vol. 4, no. 12, 2018.

[41] T. Chen, T. Lu, S. Song, S. Miao, F. Gao, and J. Li, "A deep learning method based on U-Net for quantitative photoacoustic imaging," in *Photons Plus Ultrasound: Imaging and Sensing 2020* (A. A. Oraevsky and L. V. Wang, eds.), vol. 11240, p. 112403V, International Society for Optics and Photonics, SPIE, 2020.

[42] E. D. Chan, M. M. Chan, and M. M. Chan, "Pulse oximetry: Understanding its basic principles facilitates appreciation of its limitations," *Respiratory Medicine*, vol. 107, no. 6, pp. 789–799, 2013.

[43] H. Dehghani, M. E. Eames, P. K. Yalavarthy, S. C. Davis, S. Srinivasan, C. M. Carpenter, B. W. Pogue, and K. D. Paulsen, "Near infrared optical tomography using nirfast: Algorithm for numerical model and image reconstruction," *Communications in Numerical Methods in Engineering*, vol. 25, no. 6, pp. 711–732, 2009.

[44] J. R. Rajian, P. L. Carson, and X. Wang, "Quantitative photoacoustic measurement of tissue optical absorption spectrum aided by an optical contrast agent," *Optics Express*, vol. 17, no. 6, pp. 4879–4889, 2009.

[45] H. Assi, R. Cao, M. Castelino, B. Cox, F. J. Gilbert, J. Gröhl, K. Gurusamy, L. Hacker, A. M. Ivory, J. Joseph, F. Knieling, M. J. Leahy, L. Lilaj, S. Manohar, I. Meglinski, C. Moran, A. Murray, A. A. Oraevsky, M. D. Pagel, M. Pramanik, J. Raymond, M. K. A. Singh, W. C. Vogt, L. Wang, S. Yang, Members of IPASC, and S. E. Bohndiek, "A review of a strategic roadmapping exercise to advance clinical translation of photoacoustic imaging: From current barriers to future adoption," *Photoacoustics*, vol. 32, p. 100539, 2023.

[46] M. Li, Y. Tang, and J. Yao, "Photoacoustic tomography of blood oxygenation: A mini review," *Photoacoustics*, vol. 10, pp. 65–73, 2018.

[47] K. Maslov, H. F. Zhang, and L. V. Wang, "Effects of wavelength-dependent fluence attenuation on the noninvasive photoacoustic imaging of hemoglobin oxygen saturation in subcutaneous vasculature in vivo," *Inverse Problems*, vol. 23, no. 6, p. S113, 2007.

[48] H. Dehghani, M. E. Eames, P. K. Yalavarthy, S. C. Davis, S. Srinivasan, C. M. Carpenter, B. W. Pogue, and K. D. Paulsen, "Near infrared optical tomography using nirfast: Algorithm for numerical model and image reconstruction," *Communications in Numerical Methods in Engineering*, vol. 25, no. 6, pp. 711–732, 2009.

[49] N. Bosschaart, G. J. Edelman, M. C. G. Aalders, T. G. van Leeuwen, and D. J. Faber, "A literature review and novel theoretical approach on the optical properties of whole blood," *Lasers in Medical Science*, vol. 29, no. 2, p. 453–479, 2014.

[50] N. Bosschaart, G. J. Edelman, M. C. G. Aalders, T. G. van Leeuwen, and D. J. Faber, "A literature review and novel theoretical approach on the optical properties of whole blood," *Lasers in Medical Science*, vol. 29, no. 2, p. 453–479, 2014.

[51] P. Lanka, K. J. Francis, H. Kruit, A. Farina, R. Cubeddu, S. K. V. Sekar, S. Manohar, and A. Pifferi, "Optical signatures of radiofrequency ablation in biological tissues," *Scientific Reports*, vol. 11, no. 1, p. 6579, 2021.

[52] S. K. V. Sekar, I. Bargigia, A. D. Mora, P. Taroni, A. Ruggeri, A. Tosi, A. Pifferi, and A. Farina, "Diffuse optical characterization of collagen absorption from 500 to 1700 nm," *Journal of Biomedical Optics*, vol. 22, no. 1, p. 015006, 2017.

[53] S. C. Davis, H. Dehghani, J. Wang, S. Jiang, B. W. Pogue, and K. D. Paulsen, "Image-guided diffuse optical fluorescence tomography implemented with laplacian-type regularization," *Optics Express*, vol. 15, no. 7, pp. 4066–4082, 2007.

[54] N. Thuerey, P. Holl, M. Mueller, P. Schnell, F. Trost, and K. Um, *Physics-based Deep Learning*. 2021.